

BGP EVPN VXLAN Configuration Guide

First Published: 2026-09-28

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
http://www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED “AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: https://www.cisco.com/c/en/us/about/legal/trademarks.html. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2026 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1 **BGP EVPN VXLAN Overview** 1

- Restrictions for BGP EVPN VXLAN 1
- BGP EVPN VXLAN 1
- The Evolution of BGP EVPN VXLAN 1
- Benefits of Deploying Overlay-Underlay Architecture using BGP EVPN VXLAN 2
- Fundamental Concepts of BGP EVPN VXLAN 2
  - VXLAN Overlay 3
    - Virtual Tunnel End Points 3
    - Overlay Multicast 4
  - Underlay 4
  - EVPN Control Plane 5
    - Route Target 6
    - EVPN Route Types 6
    - EVPN Instance 6
    - Ethernet Segment 6
    - EVPN Multihoming 6
  - Stretched VLAN and Subnet 7
  - Spine Leaf Architecture 7
  - Integrated Routing and Bridging 9
  - VXLAN Gateways 9
  - Layer 2 Virtual Network Instance 10
  - Layer 3 Virtual Network Instance 10
  - Mobility 10

CHAPTER 2 **Configuring EVPN VXLAN Layer 2 Overlay Network** 13

- Restrictions for Configuring EVPN VXLAN Layer 2 Overlay Network 13

Information About EVPN VXLAN Layer 2 Overlay Network 13
Broadcast, Unknown Unicast, and Multicast Traffic 13
Underlay Multicast 14
Ingress Replication 14
BUM Traffic Rate Limiting 15
Flooding Suppression 15
How to Configure EVPN VXLAN Layer 2 Overlay Network 16
Configuring Layer 2 VPN EVPN on a VTEP 17
Configuring an EVPN Instance on the VLAN on a VTEP 20
Configuring the Access-Facing Interface in the VLAN on a VTEP 21
Configuring the Loopback Interface on a VTEP 21
Configuring the NVE Interface on a VTEP 22
Configuring BGP on a VTEP with EVPN Address Family 23
Verifying EVPN VXLAN Layer 2 Overlay Network 25
Configuration Examples for EVPN VXLAN Layer 2 Overlay Network 27
Example: Configuring Layer 2 VNI with Back-to-Back Multicast Replication 27
Example: Configuring Layer 2 VNI with Back to Back Ingress Replication 35
Example: Configuring Layer 2 VNI with Spine Multicast Replication 41
Example: Configuring Layer 2 VNI with Spine Ingress Replication 59

CHAPTER 3 **BGP EVPN Layer 2 Overlay CLI Simplification 73**
Simplified CLIs for Layer 2 Overlay Configuration 73
Service Profile 73

CHAPTER 4 **Configuring EVPN VXLAN Layer 2 Overlay with Q-in-VNI 77**
Restrictions for EVPN VXLAN Layer 2 Overlay with Q-in-VNI 77
Information About Q-in-VNI 77
How to Configure Q-in-VNI in a BGP EVPN VXLAN Fabric 79
Example: Configuring Q-in-VNI in a BGP EVPN VXLAN Fabric 80
Additional References for EVPN VXLAN Layer 2 Overlay with Q-in-VNI 81

CHAPTER 5 **Configuring EVPN VXLAN Layer 3 Overlay Network 83**
Information About EVPN VXLAN Layer 3 Overlay Network 83
How to Configure EVPN VXLAN Layer 3 Overlay Network 84

Configuring an IP VRF on a VTEP 84
Configuring the Core-facing VLAN on a VTEP 86
Configuring Access-facing VLAN on a VTEP 87
Configuring Switch Virtual Interface for the Core-facing VLAN 87
Configuring the Switch Virtual Interface for the Access-facing VLANs 88
Configuring the Loopback Interface on a VTEP 89
Configuring the NVE Interface on a VTEP 90
Configuring BGP with EVPN and VRF Address Families on a VTEP 91
Configuration Examples for EVPN VXLAN Layer 3 Overlay Network 94
Verifying EVPN VXLAN Layer 3 Overlay Network 101

CHAPTER 6 **Configuring EVPN VXLAN Integrated Routing and Bridging 103**
Restrictions for EVPN VXLAN Integrated Routing and Bridging 103
Integrated Routing and Bridging 104
EVPN VXLAN Distributed Anycast Gateway 104
Manual MAC Address Configuration 105
MAC Aliasing 106
EVPN VXLAN Centralized Default Gateway 106
Default Gateway MAC Address Assignment 107
Route Type 2 to Route Type 5 Host Route Reorigination for Overlay Networks with a Centralized Gateway (Asymmetric IRB) 108
How to Configure EVPN VXLAN Integrated Routing and Bridging 108
Configuring EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway 108
Configuring Layer 2 VPN EVPN on a VTEP 109
Configuring an IP VRF on a VTEP 112
Configuring the Loopback Interface on a VTEP 114
Configuring Core-facing and Access-facing VLANs on a VTEP 114
Configuring Switch Virtual Interface for the Core-facing VLAN on a VTEP 116
Configuring Switch Virtual Interface for the Access-facing VLANs on a VTEP 117
Configuring the NVE Interface on a VTEP 118
Configuring BGP with EVPN and VRF Address Families on a VTEP 119
Configuring EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway 122
Verifying EVPN VXLAN Integrated Routing and Bridging 123
Configuration Examples for EVPN VXLAN Integrated Routing and Bridging 123

Example: Enabling EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway 124
Example: Enabling EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway 138

CHAPTER 7 **Configuring DHCP Relay in a BGP EVPN VXLAN Fabric** 155
Information About DHCP Relay in a BGP EVPN VXLAN Fabric 155
DHCP Relay on VTEPs in Distributed Anycast Gateway Deployment 156
DHCP Relay on VTEPs in a Layer 2 Overlay Fabric 157
How to Configure DHCP Relay in a BGP EVPN VXLAN Fabric 157
Configuring DHCP Relay on a VTEP 157
Configuring DHCP Relay on the Access SVI of a VTEP 158
Configuring the Layer 3 or Routed Interface on the Border VTEP for DHCP Server Reachability 160
Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric 162
Example: DHCP Server is in the Layer 3 Default VRF and the DHCP Client is in the Tenant VRF 163
Example: DHCP Server and DHCP Client are in the Same Tenant VRF 167
Example: DHCP Client and DHCP Server are in Different Tenant VRFs 170
Example: DHCP Server is in a non-Default, non-VXLAN VRF and DHCP Client is in the Tenant VRF 175
Example: Layer 3 Gateway Configured on Border VTEP 180
Example: Layer 3 Gateway Configured on External Router or Firewall 181

CHAPTER 8 **Configuring VXLAN-Aware Flexible Netflow** 187
Restrictions for VXLAN-Aware Flexible NetFlow 187
Information About VXLAN-Aware Flexible NetFlow 187
How to Configure VXLAN-Aware Flexible NetFlow 187
Configuring a Flow Record 188
Configuring Flow Exporter 189
Configuring a Flow Monitor 191
Configuring Flexible NetFlow on an NVE Interface 191
Configuration Examples for VXLAN-Aware Flexible NetFlow 192

CHAPTER 9 **Configuring BGP EVPN VXLAN QoS** 199
Quality of Service Components 199

How to Configure BGP EVPN VXLAN QoS 200
BGP EVPN VXLAN QoS Modes 201
BGP EVPN VXLAN Traffic without QoS Policy 202
BGP EVPN VXLAN QoS Marking with an Input Policy 203
BGP EVPN VXLAN QoS Marking with Input and Output Policies 204

CHAPTER 10 **Configuring Optimized Layer 2 Overlay Multicast** 205
Restrictions for Optimized Layer 2 Overlay Multicast 205
Information About Optimized Layer 2 Overlay Multicast 205
Optimized Layer 2 Multicast in BGP EVPN VXLAN Layer 2 Overlay 206
Optimized Layer 2 Overlay Multicast Interworking in BGP EVPN VXLAN Fabric 207
Optimized Layer 2 Multicast Across Layer 2 Overlays 207
Optimized Layer 2 Overlay Multicast Handoff with External Network 208
Replication Types Supported in an Underlay Network 208
How to Configure Optimized Layer 2 Overlay Multicast 209
Configure Optimized Layer 2 Multicast in EVPN VXLAN Layer 2 Overlay 209
Enable Optimized Layer 2 Overlay Multicast in BGP EVPN VXLAN Fabric 209
Configure Optimized Layer 2 Overlay Multicast for IPv4 210
Configure Optimized Layer 2 Overlay Multicast for IPv6 211
Configure Optimized Layer 2 Overlay Multicast Interworking 212
Enable Layer 3 Multicast on Centralized Gateway 212
Verifying Optimized Layer 2 Overlay Multicast on VTEP and Centralized Gateway 213
Configuration Examples for Optimized Layer 2 Overlay Multicast 214
Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication 215
Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay 222
Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 230
Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic 237
Example 3: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication 257
Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay 264
Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication 274
Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic 282

Example 5: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 313

Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 321

Example 6: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6 341

Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6 349

CHAPTER 11 **Configuring Tenant Routed Multicast** 361

Restrictions for Tenant Routed Multicast 361

Information about Tenant Routed Multicast 361

TRM in PIM Sparse Mode 363

RP Placement 363

PIM Sparse Mode with Anycast RP 366

PIM Sparse Mode with RP Inside the BGP EVPN VXLAN Fabric 367

PIM Sparse Mode with RP Outside the BGP EVPN VXLAN Fabric 369

TRM in PIM Source Specific Mode 370

Data MDT 371

How to Configure Tenant Routed Multicast 373

Configuring Tenant Routed Multicast (TRM) 373

Configuring Tenant Routed Multicast (TRM) Advanced Mode 375

Configuring TRM with PIM Sparse Mode 375

Configuring TRM with PIM Source Specific Mode 382

Verifying Tenant Routed Multicast 382

Configuration Examples for Tenant Routed Multicast 383

Example: Configuring TRM in PIM Sparse Mode with Anycast RP 383

Example: Configuring TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric 423

Example: Configuring TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric 456

Example: Configuring TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric 495

Example: Configuring TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when RP is Outside the BGP EVPN VXLAN Fabric 526

Troubleshooting Tenant Routed Multicast 566

CHAPTER 12 **Configuring Multi-Homing in a BGP EVPN VXLAN Fabric** 567

Restrictions for Multi-Homing in a BGP EVPN VXLAN Fabric 567

Information About Multi-Homing in a BGP EVPN VXLAN Fabric 567

Single-Homing 568

Multi-Homing 569

Single-Active Redundancy Mode 569

Migration Between Single-Homed and Multi-Homed Network Topologies 572

How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric 572

Configuring Dual-Homing with Single Active Redundancy in a BGP EVPN VXLAN Fabric 572

Configuring Redundancy on an Ethernet Segment 572

Associating an Ethernet Segment with an Interface on a VTEP 573

Migrating from a Single-Homed Topology to a Single-Active Dual-Homed Topology 574

Migrating from a Single-Active Dual-Homed Topology to a Single-Homed Topology 576

Configuration Examples for Multi-Homing in a BGP EVPN VXLAN Fabric 577

Example: Configuring Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric 577

CHAPTER 13 **Configuring Private VLANs in a BGP EVPN VXLAN Fabric** 599

Restrictions for Private VLANs in a BGP EVPN VXLAN Fabric 599

Information About Private VLANs in a BGP EVPN VXLAN Fabric 599

Primary and Secondary VLANs 600

Private VLAN Ports 600

Extension of Private VLANs in a BGP EVPN VXLAN Fabric 600

Traffic Forwarding for Private VLANs in a BGP EVPN VXLAN Fabric 601

Known Unicast Traffic Forwarding 601

Broadcast, Unknown Unicast, and Multicast Traffic Forwarding 603

Routed Traffic Forwarding 605

How to Configure Private VLANs in a BGP EVPN VXLAN Fabric 605

Configuring the Primary and Secondary VLANs for a Private VLAN 606

Configuring the Port of a Private VLAN 607

Enabling EVPN in a Private VLAN 608

Configuration Examples for Private VLANs in a BGP EVPN VXLAN Fabric 608

CHAPTER 14 **Configuring EVPN Microsegmentation** 633

Restrictions for EVPN Microsegmentation 633

Information About EVPN Microsegmentation 633

Classification 634

Propagation 634

Enforcement 635

Security Group Access Control Lists 635

How to Configure EVPN Microsegmentation 636

Deployment of EVPN Microsegmentation 636

Leaf Node at Access or Distribution Layer 638

Device Tracking for a Trunk Interface 639

Configuration Example of EVPN Microsegmentation 640

CHAPTER 15 **Configuring Additional BGP EVPN Features** 641

Dynamic BGP Peering for EVPN 641

Configuration Example for Dynamic BGP Peering 641

EVPN Route Map Support 641

Configuring Route Maps 642

Matching based on IP prefix 642

Matching Based on EVPN Route Type 643

Applying Route Maps 644

BGP EVPN VRF Auto Route Distinguisher 645

Prerequisites for Configuring Auto Route Distinguisher in a BGP EVPN VXLAN Fabric 645

Configuring Auto Route Distinguisher in a BGP EVPN VXLAN Fabric 645

BGP EVPN VRF Auto Route Target 646

Prerequisites for Configuring Auto Route Target in a BGP EVPN Fabric 647

Configuring Auto Route Target in a BGP EVPN VXLAN Fabric 647

CHAPTER 16 **Troubleshooting BGP EVPN VXLAN** 651

Troubleshooting Scenarios for BGP EVPN VXLAN 651

Troubleshooting Broadcast, Unkown Unicast, Multicast Traffic Forwarding 652

Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI 656

Verifying the Provisioning of an EVPN VXLAN Layer 2 Overlay Network 657

Verifying Intra-Subnet Traffic Movement in an EVPN VXLAN Layer 2 Overlay Network 661
Troubleshooting Unicast Forwarding Between VTEPS in Different VLANs Through a Layer 3 VNI 668
Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network 669
Verifying Inter-Subnet Traffic Movement and Symmetric IRB in an EVPN VXLAN Layer 3 Overlay Network 674
Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network 681
Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network 681
Verifying Traffic from a VXLAN Fabric to an IP Network Through a Border Leaf Switch Using Route Type 5 685
Troubleshooting Tenant Routed Multicast 688

PART I **EVPN Multihoming: All-Active Mode** 689

CHAPTER 17 **All-Active Mode** 691
EVPN multihoming key benefits 691
Restrictions for EVPN multihoming in all-active mode 692
EVPN multihoming technology overview 693
Network system components 693
EVPN multihoming control plane 694
Ethernet segment 694
Ethernet segment ID 695
Anycast gateway MAC and IP addresses 695
Designated forwarder and non-designated forwarder roles 695
Dynamic MAC and IP learning and synchronization 696
IGMP join and leave synchronization 696
EVPN multihoming for non-fabric networks 696
EVPN multihoming in fabric networks 697
Single unified BGP control plane 697
Non-hierarchical fabric networks 698
EVPN fabric overlay network types overview 698
EVPN fabric and non-fabric network co-existence 700

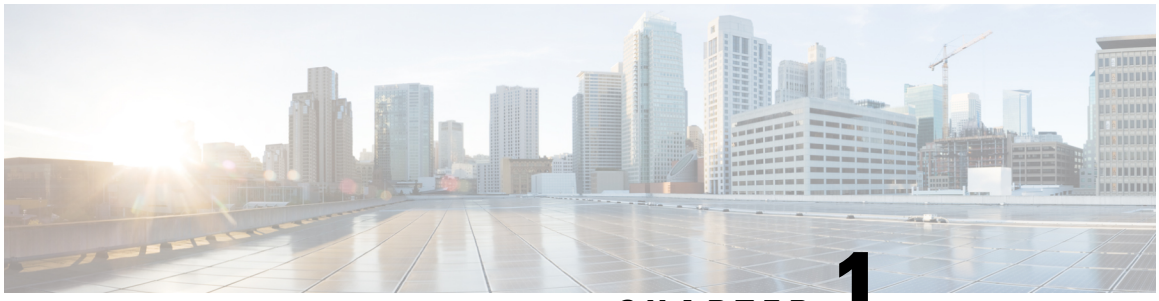
CHAPTER 18 **Overlay: Routed** 703
Routed overlay network overview 703

Understanding routed overlay networks 704
Configure routed overlay networks 705
Configure IP VRFs 706
Configure core VLAN and SVI interfaces 707
Configure access SVI interfaces 709
Configure spine EVPN route policies 710
Configure IP VRF address family routing 712
Verify the routed overlay network configuration 713
Reference routed overlay network configuration 715
Layer 2 EVPN multihoming 716
Underlay: fabric core and BGP peering 718
Overlay routed networks 723

CHAPTER 19 **Overlay: Distributed Anycast Bridged** 727
Overview: Distributed Anycast Gateway bridged overlay networks 727
Guidelines for DAG bridged overlay networks 728
Understanding DAG bridged overlay networks 728
Configure DAG bridged overlay networks 730
Configure IP VRFs 730
Configure core VLAN and SVI interfaces 732
Configure a DAG bridged overlay MAC-VRF 734
Configure access SVI interfaces 736
Configure a spine policy for DAG bridged overlay networks 737
Configure IP VRF address family routing under BGP 739
Verify the DAG bridged overlay network configuration 740
Reference configuration for DAG bridged overlay networks 743
Layer 2 EVPN multihoming 743
Underlay: fabric core and BGP peering 749
Overlay: DAG bridged networks 757

CHAPTER 20 **Overlay:Distributed Anycast Gateway Routed** 761
Distributed Anycast Gateway: routed overlay network overview 761
Restrictions for DAG routed overlay networks 762
Understanding DAG routed overlay networks 762

Configure DAG routed overlay networks 764
Configure IP VRFs 764
Configure core VLAN and SVI interfaces 766
Configure DAG routed MAC VRFs 768
Configure IPv4 ARP and IPv6 ND/DAD proxy 769
Configure spine policies for DAG routed overlay networks 771
Configure IP VRF address family routing under BGP 772
Verify the DAG routed overlay configuration 773
Reference configuration for DAG routed overlay networks 776
Layer 2 EVPN multihoming 777
Underlay: fabric core and BGP peering 783
Overlay: DAG routed networks 790



CHAPTER 1

BGP EVPN VXLAN Overview

- Restrictions for BGP EVPN VXLAN, on page 1
- BGP EVPN VXLAN, on page 1
- The Evolution of BGP EVPN VXLAN, on page 1
- Benefits of Deploying Overlay-Underlay Architecture using BGP EVPN VXLAN, on page 2
- Fundamental Concepts of BGP EVPN VXLAN, on page 2

Restrictions for BGP EVPN VXLAN

BGP EVPN VXLAN does not support Generic Routing Encapsulation (GRE) in the underlay.

BGP EVPN VXLAN

BGP EVPN VXLAN is a campus network solution for Cisco Industrial Ethernet switches running Cisco IOS XE software. This solution is a result of ratified IETF RFC specifications RFC 7432 and RFC 8365 supporting BGP EVPN control plane with RFC 7348 VXLAN data-plane. In addition, multiple other BGP EVPN RFCs and Internet drafts submitted by the BGP Enabled ServicesS (bess[1]) workgroup are supported. It is designed to provide a unified overlay network solution and also address the challenges and drawbacks of existing technologies.

This chapter provides a background for the evolution of the solution and covers conceptual information and basic terminology that is required to understand BGP EVPN VXLAN. Later chapters of this configuration guide cover information about configuration, implementation, functionalities, and troubleshooting BGP EVPN VXLAN.

The Evolution of BGP EVPN VXLAN

Traditionally, VLANs have been the standard method for providing network segmentation in campus networks. VLANs use loop prevention techniques such as Spanning Tree Protocol (STP), which impose restrictions on network design and resiliency. Further, there is a limitation with the number of VLANs that can be used to address layer 2 segments (4094 VLANs). Therefore, VLANs are a limiting factor for IT departments and cloud providers who build large and complex campus networks.

VXLAN is designed to overcome the inherent limitations of VLANs and STP. It is a proposed IETF standard [RFC 7348] to provide the same Ethernet Layer 2 network services as VLANs do, but with greater flexibility.

Functionally, VXLAN is a MAC-in-UDP encapsulation protocol that runs as a virtual overlay on an existing Layer 3 network.

However, VXLAN by itself does not provide optimal switching and routing capabilities in a network because it uses a “flood and learn” mechanism that limits scalability. "Flood and learn" mechanism is where the host’s information is flooded across the network for it to be reachable. To provide optimal switching and routing capabilities, a VXLAN overlay requires:

- An underlying transport network that performs data plane forwarding to allow unicast communication between end points connected to the fabric.
- A control plane that is capable of distributing Layer 2 and Layer 3 host reachability information across the network.

To meet these additional requirements, Internet drafts submitted by the bess workgroup (draft-ietf-bess-evpn-overlay-12) proposed MP-BGP to carry Layer 2 MAC and Layer 3 IP information simultaneously. MP-BGP incorporates Network Layer Reachability Information (NLRI) to achieve this. With MAC and IP information available together for forwarding decisions, routing and switching within a network is optimized. This also minimizes the use of the conventional "flood and learn" mechanism used by VXLAN and allows for scalability in the fabric. EVPN is the extension that allows BGP to transport Layer 2 MAC and Layer 3 IP information. This deployment is called a BGP EVPN VXLAN fabric (also referred to as VXLAN fabric).

Benefits of Deploying Overlay-Underlay Architecture using BGP EVPN VXLAN

Deploying an overlay-underlay architecture using BGP EVPN VXLAN provides the following advantages:

- Scalability — VXLAN provides Layer 2 connectivity that allows for infrastructure that can scale to 16 million tenant networks. It overcomes the 4094-segment limitation of VLANs. This is necessary to address today’s multi-tenant cloud requirements.
- Flexibility — VXLAN allows workloads to be placed anywhere, along with the traffic separation required, in a multitenant environment. The traffic separation is done by network segmentation using VXLAN segment IDs or VXLAN network identifiers (VNIs). Workloads for a tenant can be distributed across different physical devices but they are identified by their respective Layer 2 VNI or Layer 3 VNI.
- Mobility — Virtual machines can be moved from one location to another without updating spine switch tables. This is because entities within the same tenant VXLAN network retain the same VXLAN segment ID, regardless of their location.

Fundamental Concepts of BGP EVPN VXLAN

This section provides information about the various fundamental concepts and terminologies that are involved in the working of BGP EVPN VXLAN.

VXLAN Overlay

An overlay network is a virtual network that is built over an existing Layer 2 or Layer 3 network by forming a static or dynamic tunnel that runs on top of the physical network infrastructure. The existing Layer 2 or Layer 3 network is what forms the underlay and is covered further below in this chapter.

When a data packet is sent through an overlay, the original packet or frame is packaged or encapsulated at a source edge device with an outer header and dispatched toward an appropriate destination edge device. The intermediate network devices forward the packet based on the outer header but are not aware of the data in the original packet. At the destination edge device, the packet is decapsulated by stripping off the overlay header and then forwarded based on the actual data within.

In the context of BGP EVPN VXLAN, VXLAN is used as the overlay technology to encapsulate the data packets and tunnel the traffic over a Layer 3 network. VXLAN creates a Layer 2 overlay network by using a MAC-in-UDP encapsulation. A VXLAN header is added to the original Layer 2 frame and it is then placed within a UDP-IP packet. A VXLAN overlay network is also called as a VXLAN segment. Only host devices and virtual machines within the same VXLAN segment can communicate with each other.

VXLAN Network Identifier

Each VXLAN segment is identified through a 24-bit segment ID, termed the VXLAN network identifier. This ensures that up to 16 million VXLAN segments can be present within the same administrative domain.

Virtual Tunnel End Points

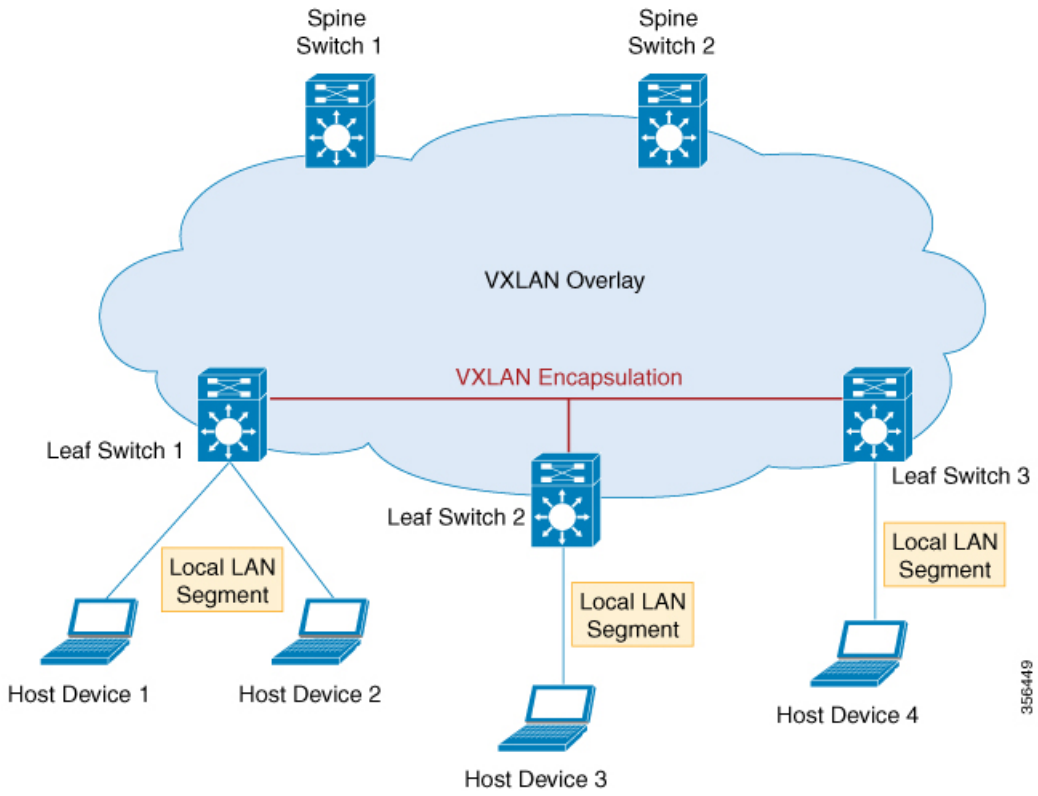
Every VXLAN segment has tunnel edge devices known as Virtual Tunnel End points (VTEPs). These devices sit at the edge of the VXLAN network and are responsible for creating instances of VXLAN tunnels, and for performing VXLAN encapsulation and decapsulation.

A VTEP has a switch interface on the local LAN segment to support local endpoint communication through bridging, and an IP interface to interact with the transport IP network.

The IP interface has a unique IP address that identifies the VTEP on the transport IP network. The VTEP uses this IP address to encapsulate Ethernet frames and transmits the encapsulated packets to the transport network through the IP interface. A VTEP device also discovers the remote VTEPs for its VXLAN segments and learns remote MAC address-to-VTEP mappings through its IP interface.

The following figure illustrates the working of an overlay VXLAN network connecting various VTEPs:

Figure 1: Overlay VXLAN Network



Overlay Multicast

Overlay multicast is the method by which a overlay network forwards multicast traffic between various VTEPs present in the network. Tenant Routed Multicast (TRM) provides a mechanism to efficiently forward multicast traffic in a VXLAN overlay network. TRM is a BGP-EVPN based solution that enables multicast routing between sources and receivers connected on VTEPs in VXLAN fabric.

Without TRM, the multicast traffic is sent as part of the underlay network in the form of BUM traffic either using underlay multicast or ingress replication methods. This does not allow sources and receivers that are present across different subnets to communicate with each other. Using TRM, multicast communication is moved out of the BUM underlay traffic. This enables multicast communication in the overlay network irrespective of the subnet in which the source or the receiver resides.

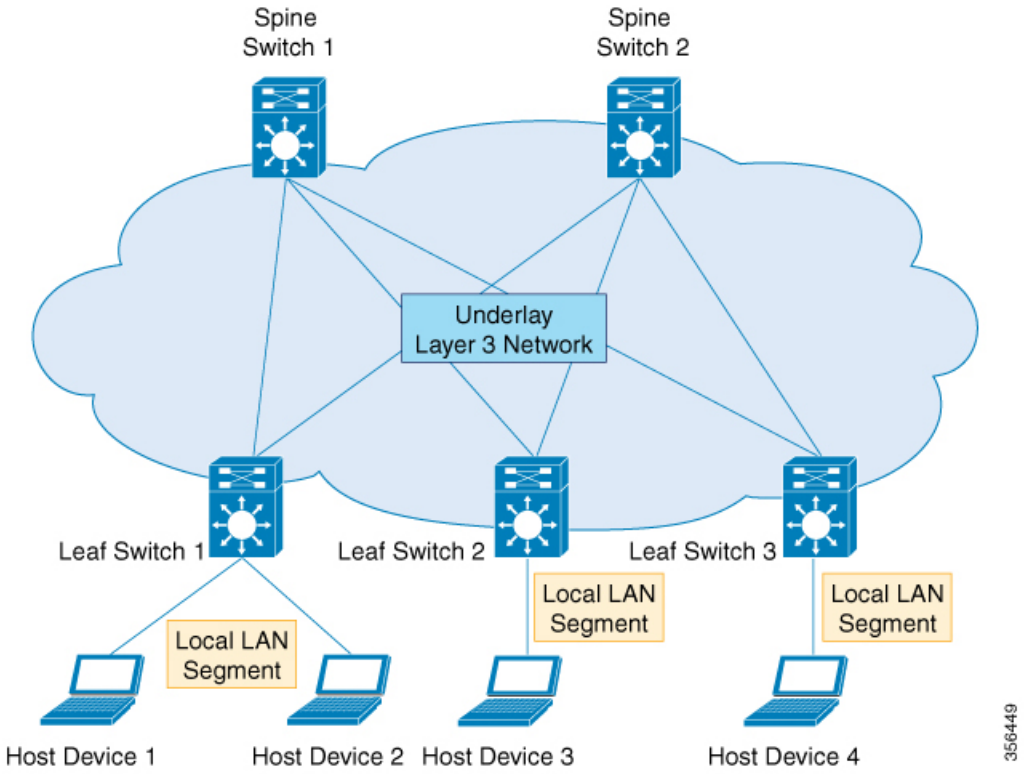
Underlay

An underlay network is the physical network over which the virtual overlay network is established. Once the overlay network is defined along with the data-plane encapsulation, a method of transport is required to move the data across the physical network underneath. This method of transport is typically an underlay transport network, or simply the underlay.

In BGP EVPN VXLAN, the underlay Layer 3 network transports the VXLAN-encapsulated packets between the source and destination VTEPs and provides reachability between them. The VXLAN overlay and the underlying IP network between the VTEPs are independent of each other.

The following figure illustrates an underlay network:

Figure 2: Underlay Layer 3 Network in BGP EVPN VXLAN



EVPN Control Plane

The overlay requires a mechanism to know which end host device is behind which overlay edge device. VXLAN natively operates on a flood and learn mechanism where broadcast, unknown unicast and multicast (BUM) traffic in a given VXLAN network is sent over the IP core to every VTEP that has membership in that network. IP multicast is used to send traffic over the network. The receiving VTEPs decapsulate the packet and, based on the inner frame, perform Layer 2 MAC learning. The inner source MAC address is learned against the outer source IP address corresponding to the source VTEP. In this way, reverse traffic is unicasted toward the previously learnt end host.

The drawback of the flood and learn mechanism is that it does not allow scalability in a VXLAN network. In order to address this issue, a control plane is used to manage the MAC address learning and VTEP discovery. In BGP EVPN VXLAN deployments, Ethernet Virtual Private Network (EVPN) is used as the control plane. EVPN control plane provides the capability to exchange both MAC address and IP address information. EVPN uses Multi Protocol Border Gateway Protocol (MP-BGP) as the routing protocol to distribute reachability information pertaining to the VXLAN overlay network, including endpoint MAC addresses, endpoint IP addresses, and subnet reachability information. BGP EVPN distribution protocol facilitates the mapping information to be built by the tunnel edge devices in the location-identity mapping database.

Route Target

A route target is an extended attribute in EVPN route updates that controls route distribution in a multi-tenant network. EVPN VTEPs have an import route target setting and an export route target setting for every VRF and Layer 2 Virtual Network Instance (VNI). When a VTEP advertises EVPN routes, it affixes its export route target in the route update. These routes are received by the other VTEPs in the network. The receiving VTEPs compare the route target value carried with the route against their own local import route target setting. If the two values match, the route is accepted and programmed in the routing table. Otherwise, the route is not imported.

EVPN Route Types

The EVPN control plane advertises the following types of information:

- Route type 1 – This is an Ethernet Auto-Discovery (EAD) route type used to advertise Ethernet segment identifier, Ethernet Tag ID, and EVPN instance information. EAD route advertisements may be sent for each EVPN instance or for each Ethernet segment.
- Route type 2 – This advertises endpoint reachability information, including MAC and IP addresses of the endpoints or VTEPs.
- Route type 3 – This performs multicast router advertisement, announcing the capability and intention to use ingress replication for specific VNIs.
- Route type 4 – This is an Ethernet Segment route used to advertise the Ethernet segment identifier, IP address length, and the originating router's IP address.
- Route type 5 – This is an IP prefix route used to advertise internal IP subnet and externally learned routes to a VXLAN network.

EVPN Instance

An EVPN Instance (EVI) represents a Virtual Private Network (VPN) on a VTEP. It is the equivalent of IP VRF in Layer 3 VPN and is also known as a MAC VRF.

Ethernet Segment

An Ethernet segment is associated with an access-facing interface of a VTEP and represents the connection with a host device. Each Ethernet segment is assigned a unique value known as Ethernet segment identifier (ESI). When a host device is connected to more than one VTEPs, then the ESI for these connections remains the same.

EVPN Multihoming

EVPN multihoming allows you to connect a Layer 2 device or an end host device to more than one leaf switch in the VXLAN network. This provides redundancy and allows network optimization over single-homed topologies where the customer network is connected to a single leaf switch. Redundancy in the connection with the leaf switches ensures that there is no traffic disruption when there is a network failure. Multihomed topologies are more resilient, secure and efficient than single-homed topologies. EVPN multihoming operates in single-active and all-active redundancy modes.

Stretched VLAN and Subnet

By running over the existing networking infrastructure, EVPN VXLAN provides a means to stretch a Layer 2 network. EVPN VXLAN overlay allows Layer 2 segments and broadcast domains to be extended across sites or campus buildings over a Layer 3 core network. Layer 2 extension with EVPN VXLAN simplifies end user IP address management and provides seamless mobility in large campus networks.

Spine Leaf Architecture

Spine-leaf architecture is a two-layer network topology where one layer is composed of leaf switches and the other layer has one or more spine switches. This design connects all the leaf switches by providing multiple paths through the various spine switches.

Spine Switch

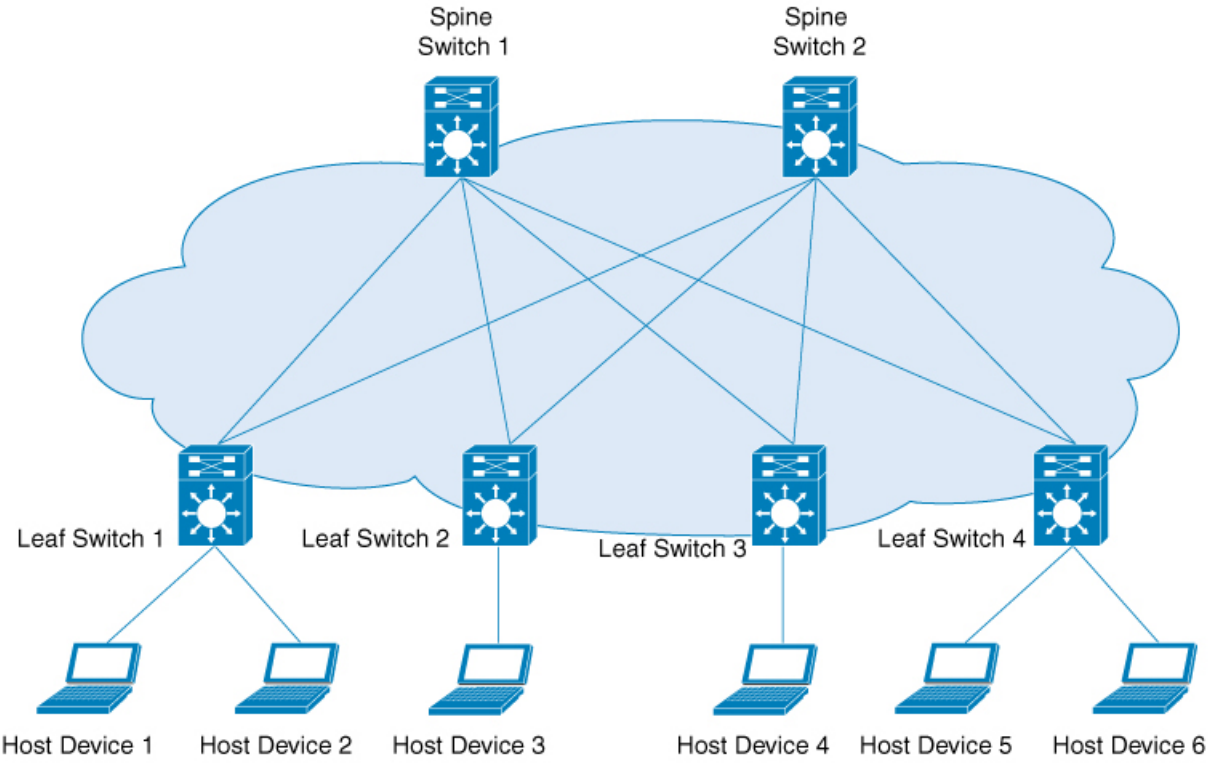
Spine switches are the connecting nodes between all the leaf switches. They forward the traffic between the leaf switches and are unaware of the endpoint addresses. By providing multiple paths to connect the leaf switches, spine switches provide redundancy to the network.

Leaf Switch

Leaf switches are the nodes that are connected to the host or access devices. As a leaf switch sits on the edge of the network, it is also called as an edge or Network Virtualization Edge (NVE). When a host device on one leaf switch tries to communicate with a host device on another leaf switch, the traffic between the leaf switches is sent through a spine switch. Leaf switches function as VTEPs in a VXLAN network and perform the encapsulation and decapsulation.

The following image shows a typical spine-leaf topology where four leaf switches are connected through two spine switches:

Figure 3: Spine-Leaf Topology

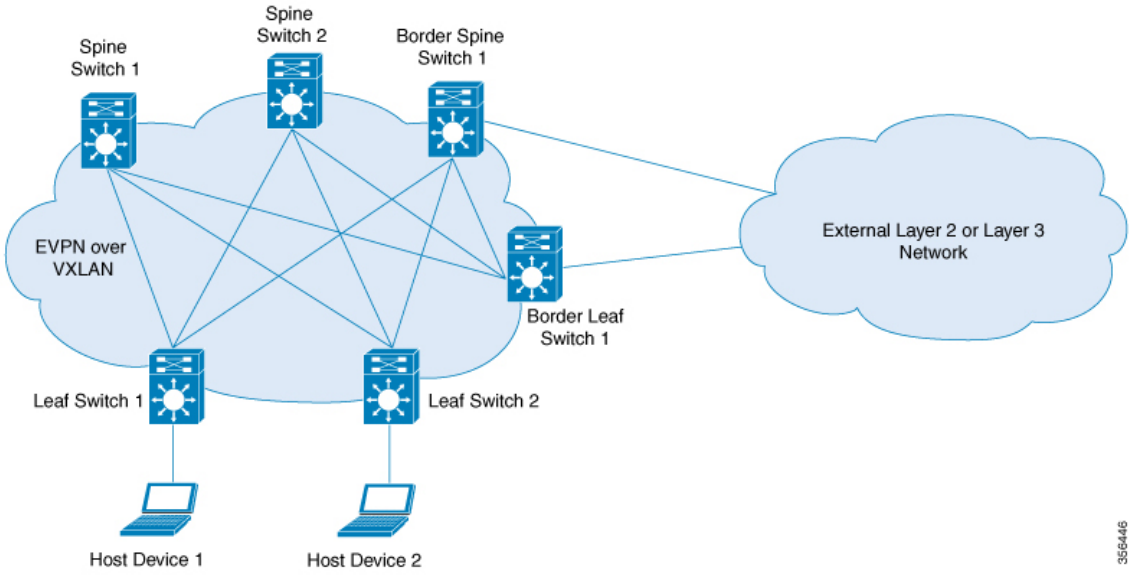


Border Spine Switch and Border Leaf Switch

External connectivity of the VXLAN fabric with other Layer 2 and Layer 3 networks is facilitated through nodes known as border nodes. If the border functionality is established through a spine switch, it is known as a border spine switch. If it is established through a leaf switch, it is known as a border leaf switch.

The following image shows a spine-leaf topology with one border leaf switch and one border spine switch connecting the fabric with an external network:

Figure 4: Spine-Leaf Topology Connected to External Network



Integrated Routing and Bridging

EVPN VXLAN supports Integrated Routing and Bridging (IRB) functionality which allows the VTEPs in a VXLAN network to forward both Layer 2 (bridged) and Layer 3 (routed) traffic. When a VTEP forwards Layer 2 traffic, it is said to be performing bridging. Similarly, when a VTEP forwards Layer 3 traffic, it is said to be performing routing. The traffic between different subnets is forwarded through the VXLAN gateways. IRB is implemented in two ways:

- Asymmetric IRB
- Symmetric IRB

VXLAN Gateways

A VXLAN Gateway is an entity in the network that forwards traffic between VXLAN segments, or from a VXLAN environment to a non-VXLAN environment. Leaf switches in a VXLAN network can function as both Layer 2 and Layer 3 VXLAN gateways.

Layer 2 VXLAN gateways forward traffic within the same VLAN. Layer 2 VXLAN gateways allow VXLAN to VLAN bridging by mapping a VNI segment to a VLAN.

Layer 3 VXLAN gateways forward traffic to a different VLAN. Layer 3 VXLAN gateways allow both VXLAN to VXLAN routing as well as VXLAN to VLAN routing. VXLAN to VXLAN routing provides Layer 3 connectivity between two VNIs where as VXLAN to VLAN routing provides connectivity between a VNI and a VLAN.

Layer 2 Virtual Network Instance

The creation of a VXLAN overlay network allows host devices connected to various leaf nodes, that are separated by multiple Layer 3 networks, to interact as if they were connected to a single Layer 2 network, which is the VXLAN segment. This logical Layer 2 segment is called as Layer 2 VNI. The traffic that flows through a Layer 2 VNI between two VLANs within the same subnet is known as bridged traffic.

A VLAN that is locally defined on a VTEP can be mapped to a Layer 2 VNI. In order to allow host devices to connect to a Layer 2 VNI, the connected VLAN must be mapped to the Layer 2 VNI, and then the Layer 2 VNI is associated with the Network Virtualization Edge (NVE) logical interface on the VTEP.

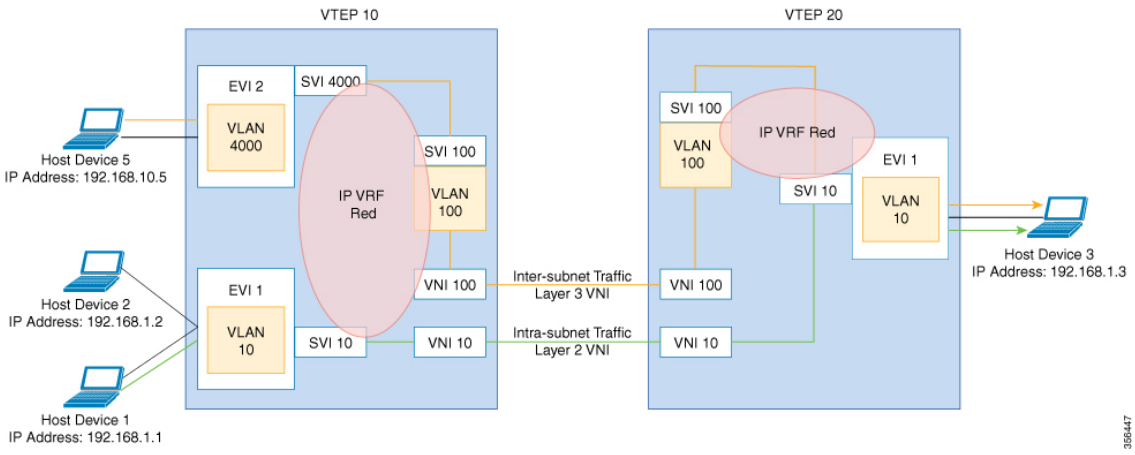
Layer 3 Virtual Network Instance

When endpoints connected to a Layer 2 VNI need to communicate with endpoints belonging to different IP subnets, they send the traffic to their default gateway. Communication between endpoints belonging to different Layer 2 VNIs is possible only through a Layer 3 routing function. In an EVPN VXLAN deployment, the various Layer 2 segments that are defined by combining the local VLANs and the global Layer 2 VNIs can be associated to a VRF in order to communicate.

A Layer 3 VNI facilitates Layer 3 segmentation for every VRF on a VTEP. This is done by mapping each VRF instance to a unique Layer 3 VNI in the network and associating the various Layer 2 VNIs for a VTEP to the same VRF. This allows inter- VXLAN communication throughout the Layer 3 VNI within a particular VRF instance. The use of VRFs to enable a logical Layer 3 isolation is known as multi-tenancy. The traffic that flows through a Layer 3 VNI between two VLANs in different subnets is known as routed traffic.

The following image shows the movement of traffic between host devices in same and different subnets through Layer 2 and Layer 3 VNIs:

Figure 5: Traffic through Layer 2 and Layer 3 VNIs



Mobility

The identity of an endpoint in the BGP EVPN control plane is derived from its MAC address and IP address, and BGP EVPN provides a mechanism to support endpoint mobility within a VXLAN overlay.

Cisco Industrial Ethernet switches support RFC 7432 which defines the scope of endpoint mobility within the VXLAN fabric.

MAC Mobility and Duplicate MAC Detection

A MAC move occurs when an endpoint (or host) moves from one port to another. The new port may be within the same VTEP, or in a different VTEP, in the same VLAN. The BGP EVPN control plane resolves such moves by advertising MAC routes (EVPN route type 2). When an endpoint's MAC address is learned on a new port, the new VTEP it is in advertises (on the BGP EVPN control plane) that it is the local VTEP for the host. All other VTEPs receive the new MAC route.

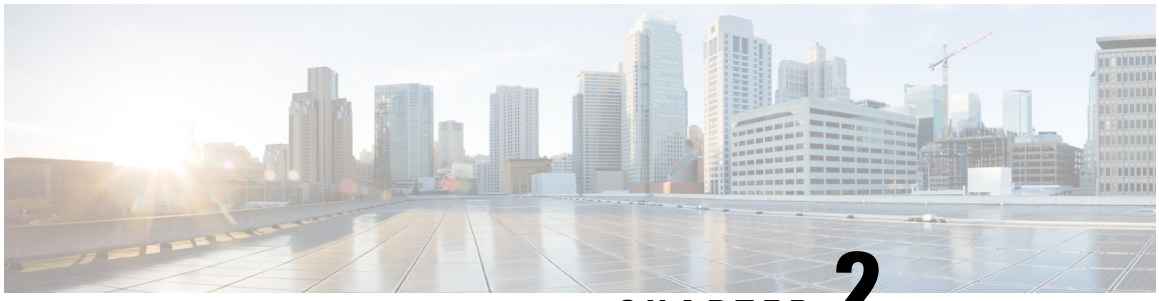
A host may move several times, causing the corresponding VTEPs to advertise as many MAC routes. There may also be a delay between the time a new MAC route is advertised and when the old route is withdrawn from the route tables of other VTEPs, resulting in two locations briefly having the same MAC route. Here, a MAC mobility sequence number helps decide the most current of the MAC routes.

When the host MAC address is learned for the first time, the MAC mobility sequence number is set to 0. The value 0 indicates that the MAC address has not had a mobility event, and the host is still at the original location. If a MAC mobility event is detected, a new Route type 2 (MAC or IP advertisement) is added to the BGP EVPN control plane by the new VTEP below which the endpoint moved (its new location). Every time the host moves, the VTEP that detects its new location increments the sequence number by 1 and then advertises the MAC route for that host on the BGP EVPN control plane. On receiving the MAC route at the old location (VTEP), the old VTEP withdraws the old route.

A case may arise in which the same MAC address is simultaneously learned on two different ports. The EVPN control plane detects this condition and alerts the user that there is a duplicate MAC. The duplicate MAC condition may be cleared either by manual intervention, or automatically when the MAC address ages out on one of the ports.

IP Mobility and Duplicate IP Detection

BGP EVPN supports IP mobility in a similar manner to the way it supports MAC mobility. The principal difference is that an IP move is detected when the IP address is learned on a different MAC address, regardless of whether it was learned on the same port or a different port. A duplicate IP address is detected when the same IP address is simultaneously learned on two different MAC addresses, and the user is alerted when this occurs.



CHAPTER 2

Configuring EVPN VXLAN Layer 2 Overlay Network

- Restrictions for Configuring EVPN VXLAN Layer 2 Overlay Network, on page 13
- Information About EVPN VXLAN Layer 2 Overlay Network, on page 13
- How to Configure EVPN VXLAN Layer 2 Overlay Network, on page 16
- Verifying EVPN VXLAN Layer 2 Overlay Network, on page 25
- Configuration Examples for EVPN VXLAN Layer 2 Overlay Network, on page 27

Restrictions for Configuring EVPN VXLAN Layer 2 Overlay Network

- Layer 2 Broadcast, Unknown Unicast, and Multicast (BUM) Traffic Forwarding using Ingress Replication.
- BUM Traffic Rate Limiting
- Dynamic ARP inspection (DAI) and DHCP Rogue Server Protection

Information About EVPN VXLAN Layer 2 Overlay Network

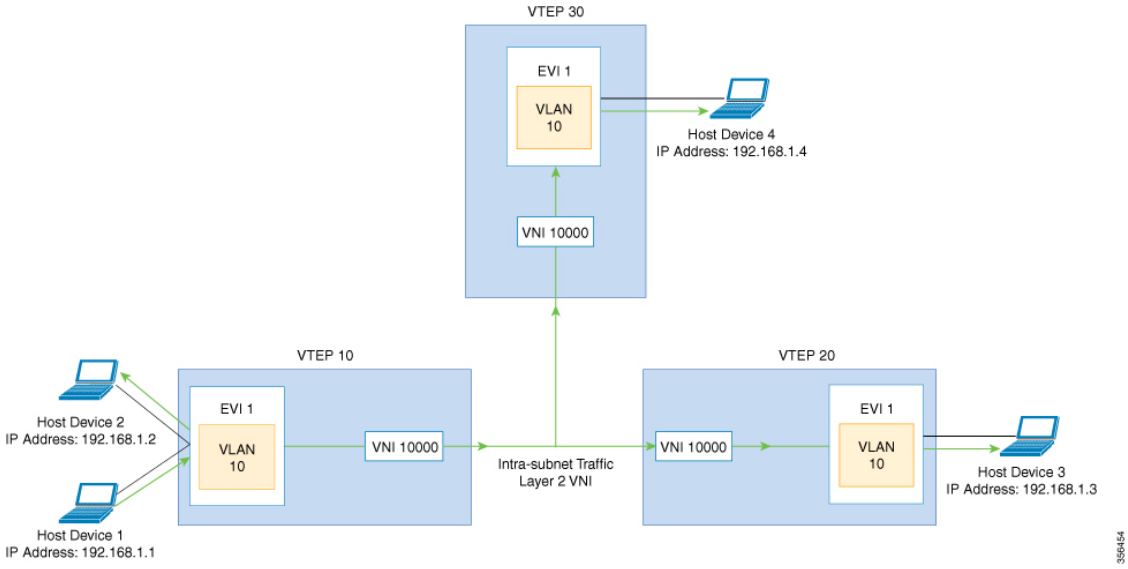
An EVPN VXLAN Layer 2 overlay network allows host devices in the same subnet to send bridged or Layer 2 traffic to each other. The network forwards the bridged traffic using a Layer 2 virtual network instance (VNI).

Broadcast, Unknown Unicast, and Multicast Traffic

Multidestination Layer 2 traffic in a VXLAN network is typically referred to as broadcast, unknown unicast, and multicast (BUM) traffic. In a BGP EVPN VXLAN fabric, the underlay network forwards the BUM traffic to all the endpoints connected to a common Layer 2 broadcast domain in the VXLAN overlay.

The following image shows the flow of BUM traffic through a Layer 2 VNI. The network forwards BUM traffic from host device 1 to all the VTEPs which in turn send the traffic to all the host devices in the same subnet.

Figure 6: BUM Traffic through Layer 2 VNI



The MP-BGP EVPN control plane uses two different methods to forward BUM traffic in a VXLAN network:

- Underlay Multicast
- Ingress Replication

Underlay Multicast

Restriction SSM is not supported for BUM Underlay Multicast.

In underlay multicast, the underlay network replicates the traffic through a multicast group. Forwarding BUM traffic using underlay multicast requires the configuration of IP multicast in the underlay network. A single copy of the BUM traffic moves from the ingress or source VTEP towards the underlay transport network. The network forwards this copy along the multicast tree so that it reaches all egress or destination VTEPs participating in the given multicast group. Various branch points in the network replicate the copy as it travels along the multicast tree. The branch points replicate the copy only if the receivers are part of the multicast group associated with the VNI.

BUM traffic forwarding through underlay multicast is achieved by mapping a Layer 2 VNI to the multicast group. This mapping must be configured on all the VTEPs associated with the Layer 2 VNI. When a VTEP joins the multicast group, it receives all the traffic that is forwarded on that group. If the VTEP receives traffic in a VNI that is not associated with it, it simply drops the traffic. This approach maintains a single link within the network, thus providing an efficient way to forward BUM traffic.

Ingress Replication

Ingress replication, or headend replication, is a unicast approach to handle multidestination Layer 2 overlay BUM traffic. Ingress replication involves an ingress device replicating every incoming BUM packet and

sending them as a separate unicast to the remote egress devices. Ingress replication happens through EVPN route type 3, also called as inclusive multicast ethernet tag (IMET) route. BGP EVPN ingress replication uses IMET route for auto-discovery of remote peers in order to set up the BUM tunnels over VXLAN. Using ingress replication to handle BUM traffic can result in scaling issues as an ingress device needs to replicate the BUM traffic as many times as there are VTEPs associated with the Layer 2 VNI.

Ingress Replication Operation

IMET routes carry the remote or egress VNIs advertised from the remote peers, which can be different from the local VNI. The network creates a VXLAN tunnel adjacency when an ingress device receives IMET ingress replication routes from remote NVE peers. The tunnel adjacency is a midchain adjacency which contains IP or UDP encapsulation for the VXLAN Tunnel. If there is more than one VNI along the tunnel, then multiple VNIs share the tunnel. Ingress replication on EVPN can have multiple unicast tunnel adjacencies and different egress VNIs for each remote peer.

The network builds a flooded replication list with the routes advertised by each VTEP. The dynamic replication list stores all the remote destination peers discovered on a BGP IMET route in the same Layer 2 VNI. The replication list gets updated every time you configure the Layer 2 VNI at a remote peer. The network removes the tunnel adjacency and VXLAN encapsulation from the replication list every time a remote NVE peer withdraws the IMET ingress replication route. The network deletes the tunnel adjacency when there is no NVE peer using it.

Any BUM traffic that reaches the ingress device gets replicated after the replication list is built. The ingress device forwards the replicated traffic throughout the network to all the remote peers in the same VNI.

BUM Traffic Rate Limiting

You can use a policer to set the flood rate limit of the BUM traffic in the network to a predefined value. This prevents the flood rate from going beyond the limit and saves the network bandwidth.

To set the flood rate limit, configure a policy with a Layer 2 miss filter on the NVE interface of a VTEP. Ensure that the policy is applied on the NVE interface for egress traffic. All the Layer 2 member VNIs under this NVE share the same policy. Any new Layer 2 VNI that is added under the NVE shares this configured policy.

Flooding Suppression

EVPN allows the distribution of the binding between IPv4 or IPv6 addresses and MAC addresses among the VTEPs of the network. It distributes the MAC-IP binding among all the VTEPs that participate in the EVPN instance associated with the MAC-IP routes. The MAC address associated with the IPv4 or IPv6 addresses is locally known even though it is learned from a remote VTEP. Locally connected endpoints send an Address Resolution Protocol (ARP) or an IPv6 neighbor discovery request when they look for a remote endpoint. The MAC-IP binding distribution allows a VTEP to perform a lookup in the local cache when it receives an ARP or an IPv6 neighbor discovery request. If the MAC-IP address information for the remote end point is available, the VTEP can use this information to avoid flooding the ARP request. If the MAC or IP address information for the remote end point is not available, the request floods throughout the fabric.

Flooding suppression avoids the flooding of ARP and IPv6 neighbor discovery (ND) packets over the EVPN VXLAN network. It suppresses the flooding to both the local and remote host or access devices. The network suppresses the flooding by implementing an ARP or neighbor discovery relay. This is achieved by using the known MAC address for the specified IPv4 or IPv6 address to convert broadcast and multicast requests to

unicast requests. Flooding suppression is enabled by default on an EVPN-enabled VLAN. An EVPN VXLAN network suppresses the flooding for the following types of traffic:

ARP Flooding Suppression

VTEPs send ARP requests as broadcast packets. ARP requests represent a large percentage of Layer 2 broadcast traffic. Flooding suppression converts them to unicast packets and reduces the network flood.

IPv6 Neighbor Discovery Flooding Suppression

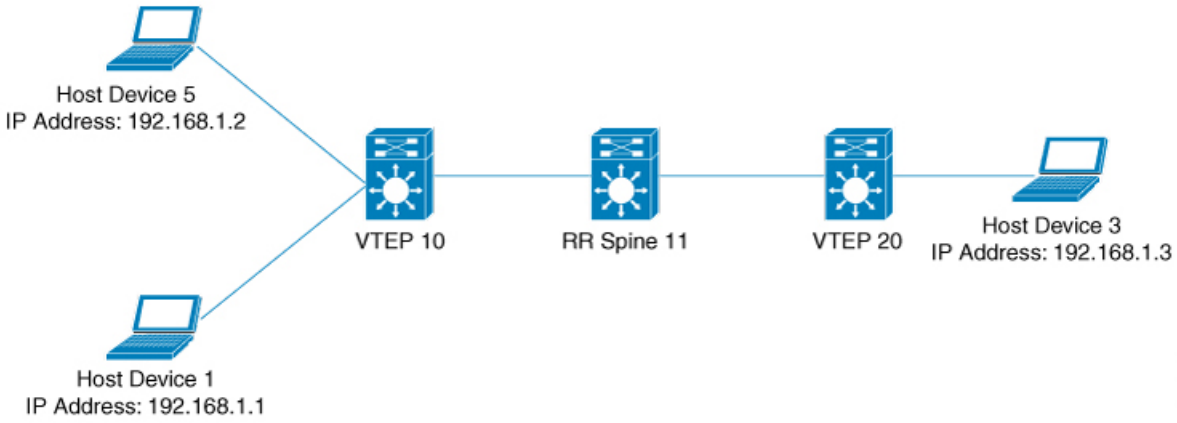
The IPv6 neighbor discovery process enables the discovery of a neighbor and helps the peers to determine each other's link-layer addresses. It also verifies the reachability of a neighbor and tracks the neighboring routers. IPv6 neighbor discovery uses Internet Control Message Protocol (ICMP) messages and solicited-node multicast addresses to achieve these functions.

Flooding suppression suppresses all multicast neighbor solicitation packets among Internet Control Message Protocol version 6 (ICMPv6) packets.

In certain scenarios where proxy ARP or proxy ND is required (for example, if the Distributed Anycast Gateway SVI interface is enabled with proxy ARP or proxy ND), use the **flooding-suppression address-resolution disable** command to explicitly disable the flooding suppression.

How to Configure EVPN VXLAN Layer 2 Overlay Network

The following figure shows a sample topology of an EVPN VXLAN Network. Host device 1 and host device 3 are part of the same subnet. The network forwards BUM traffic from host device 1 to host device 3 using a Layer 2 VNI through either underlay multicast or ingress replication methods.



Note In a two-VTEP topology, a spine switch is not mandatory. For information about configuration of spine switches in an EVPN VXLAN network, see *Configuring Spine Switches in a BGP EVPN VXLAN Fabric* module.

Perform the following set of procedures to configure an EVPN VXLAN Layer 2 overlay network and forward the BUM traffic:

- Configure Layer 2 VPN EVPN on the VTEPs.
- Configure an EVPN instance in the VLAN on the VTEPs.
- Configure the access-facing interface in the VLAN on the VTEPs.
- Configure the loopback interface on the VTEPs.
- Configure the network virtualization endpoint (NVE) interface on the VTEPs.
- Configure BGP with EVPN address family on the VTEPs.
- Configure underlay multicast, if the specified replication type is static. For more information, see *IP Multicast Routing Configuration Guide*.

Configuring Layer 2 VPN EVPN on a VTEP

To configure the Layer 2 VPN EVPN parameters on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**l2vpn evpn** **Example:** `Device(config)# l2vpn evpn`	Enters EVPN configuration mode.
Step 4	**replication-type** {**ingress** | **static**} **Example:** `Device(config-evpn)# replication-type static`	Configures the Layer 2 VPN EVPN replication type. **Note** Configure the Layer 2 VPN EVPN replication type as static, if multicast is enabled in the underlay network for EVPN BUM traffic. When the Layer 2 VPN EVPN replication type is configured as static, the IMET route is not advertised and forwarding of BUM traffic relies on underlay multicast being configured on each VTEP.

	Command or Action	Purpose
Step 5	**router-id** *loopback-interface-id* **Example:** `Device(config-evpn)# router-id loopback 0`	Specifies the interface that will supply the IP addresses to be used in auto-generating route distinguishers.
Step 6	**default-gateway advertise** **Example:** `Device(config-evpn)# default-gateway advertise`	(Optional) Enables default gateway advertisement on the switch. To configure distributed anycast gateway in a VXLAN network using MAC aliasing, enable default gateway advertisement on all the leaf switches in the network. This command is applicable in integrated routing and bridging (IRB) scenarios where Layer 2 and Layer 3 VNIs coexist in a VRF. Refer to *Configuring EVPN VXLAN Integrated Routing and Bridging* module for more details. This command is mandatory only if the same MAC address is not manually configured on all the access SVIs. **Note** Use the **default-gateway advertise** {**enable** | **disable**} command in EVPN instnace configuration mode to override the global default gateway advertisement settings and enable or disable it for a specific EVPN instance.
Step 7	**logging peer state** **Example:** `Device(config-evpn)# logging peer state`	(Optional) Displays syslog message when the first route is received or the last route is withdrawn from a given remote VTEP.
Step 8	**mac duplication limit** *limit-number* **time** *time-limit* **Example:** `Device(config-evpn)# mac duplication limit 20 time 5`	(Optional) Changes parameters for detecting duplicate MAC addresses.
Step 9	**ip duplication limit** *limit-number* **time** *time-limit* **Example:** `Device(config-evpn)# ip duplication limit 20 time 5`	(Optional) Changes parameters for detecting duplicate IP addresses.

	Command or Action	Purpose
Step 10	**route-target auto vni** **Example:** `Device(config-evpn)# route-target auto vni`	(Optional) Specifies to use VNI instead of EVPN instance number to auto-generate route target.
Step 11	**exit** **Example:** `Device(config-evpn)# exit`	Exits EVPN configuration mode and enters global configuration mode.
Step 12	**l2vpn evpn instance** *evpn-instance-number* **vlan-based** **Example:** `Device(config)# l2vpn evpn instance 1 vlan-based`	Configures a VLAN based EVPN instance in Layer 2 VPN configuration mode. An EVPN instance needs to be explicitly configured only when something needs to be configured per EVPN instance such as a route target.
Step 13	**encapsulation vxlan** **Example:** `Device(config-evpn-evi)# encapsulation vxlan`	(Optional) Defines the encapsulation format as VXLAN. The encapsulation format is VXLAN by default.
Step 14	**replication-type** {**ingress** | **static**} **Example:** `Device(config-evpn-evi)# replication-type ingress`	(Optional) Sets the replication type for the EVPN instance. In case a global replication type has already been configured, this overrides the global setting.
Step 15	**default-gateway advertise** {**enable** | **disable**} **Example:** `Device(config-evpn-evi)# default-gateway advertise disable`	(Optional) Enables or disables the default gateway advertisement for the EVPN instance. In case default gateway advertisement has already been globally configured, this overrides the global setting. This command is mandatory only if the same MAC address is not manually configured on all the access SVIs. To configure distributed anycast gateway in a VXLAN network using MAC aliasing, enable default gateway advertisement on all the leaf switches in the network.
Step 16	**ip local-learning** {**enable** | **disable**} **Example:** `Device(config-evpn-evi)# ip local-learning disable`	(Optional) Enables or disables local IP address learning for the specified EVPN instance. In case IP address learning has already been globally configured, this overrides the global setting.

	Command or Action	Purpose
Step 17	**re-originate route-type5** **Example:** `Device(config-evpn-evi)# re-originate route-type5`	(Optional) Enables the centralized gateway (CGW) VTEP to re-originate the route-type 2 (RT 2) host routes from a Layer 2 VTEP as route-type 5 (RT 5) network routes into a Layer 3 overlay network.
Step 18	**no auto-route-target** **Example:** `Device(config-evpn-evi)# no auto-route-target`	(Optional) Disables auto generation of route targets.
Step 19	**rd** *rd-value* **Example:** `Device(config-evpn-evi)# rd 65000:100`	(Optional) Configures a route distinguisher manually.
Step 20	**route-target** {**import** | **export** | **both**} *rt-value* **Example:** `Device(config-evpn-evi)# route-target both 65000:100`	(Optional) Configures route targets manually. **Note** Configure route targets manually if the auto-generated route target values (ASN:EVI or ASN:VNI) are different between the VTEPs.
Step 21	**end** **Example:** `Device(config-evpn-evi)# end`	Returns to privileged EXEC mode.

Configuring an EVPN Instance on the VLAN on a VTEP

To configure an EVPN instance on the VLAN on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan configuration** *vlan-id* **Example:** `Device(config)# vlan configuration 11`	Enters VLAN feature configuration mode for the specified VLAN interface.

	Command or Action	Purpose
Step 4	**member evpn-instance** *evpn-instance-id* **vni** *l2-vni-number* **Example:** `Device(config-vlan)# member evpn-instance 1 vni 10000`	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 2 VNI.
Step 5	**end** **Example:** `Device(config-vlan)# end`	Returns to privileged EXEC mode.

Configuring the Access-Facing Interface in the VLAN on a VTEP

To configure the access-facing interface in the VLAN on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *interface-name* **Example:** `Device(config)# interface GigabitEthernet1/0/1`	Enters interface configuration mode for the specified interface.
Step 4	**switchport access vlan** *vlan-id* **Example:** `Device(config-if)# switchport access vlan 11`	Configures the interface as a static-access port of the specified VLAN. Interface can also be configured as a trunk interface, if required.
Step 5	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring the Loopback Interface on a VTEP

To configure the loopback interface on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *loopback-interface-id* **Example:** `Device(config)# interface Loopback0`	Enters interface configuration mode for the specified Loopback interface.
Step 4	**ip address** *ipv4-address* **Example:** `Device(config-if)# ip address 10.12.11.11`	Configures the IP address for the Loopback interface.
Step 5	**ip pim sparse mode** **Example:** `Device(config-if)# ip pim sparse mode`	Enables Protocol Independent Multicast (PIM) sparse mode on the Loopback interface.
Step 6	**end** **Example:** `Device(config-vlan)# end`	Returns to privileged EXEC mode.

Configuring the NVE Interface on a VTEP

To add a VNI member to the NVE interface of a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**interface** *nve-interface-id* **Example:** `Device(config)# interface nve1`	Defines the interface to be configured as a trunk, and enters interface configuration mode.
Step 4	**no ip address** **Example:** `Device(config-if)# no ip address`	Disables IP processing on the interface by removing its IP address.
Step 5	**source-interface** *loopback-interface-id* **Example:** `Device(config-if)# source-interface loopback0`	Sets the IP address of the specified loopback interface as the source IP address.
Step 6	**host-reachability protocol bgp** **Example:** `Device(config-if)# host-reachability protocol bgp`	Configures BGP as the host-reachability protocol on the interface. **Note** You must configure the host-reachability protocol on the interface. If you do not execute this step, the VXLAN tunnel defaults to static VXLAN tunnel, which is currently not supported on the Cisco Industrial Ethernet switches.
Step 7	**member vni** *layer2-vni-id* {**ingress-replication** [**local-routing**] | **mcast-group** *multicast-group-address* **Example:** `Device(config-if)# member vni 10000 mcast-group 227.0.0.1`	Associates the Layer 2 VNI member with the NVE. The specified replication type must match the replication type that is configured globally or for the specific EVPN instance. Use **mcast-group** keyword for static replication and **ingress-replication** keyword for ingress replication. Use the **local-routing** keyword only when you need to configure route type 2 (RT 2) to route type 5 (RT 5) reorigination on the centralized gateway (CGW) VTEP.
Step 8	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring BGP on a VTEP with EVPN Address Family

To configure BGP on a VTEP with EVPN address family and with spine switch as the neighbor, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 1`	Enables a BGP routing process, assigns it an autonomous system number, and enters router configuration mode.
Step 4	**bgp log-neighbor-changes** **Example:** `Device(config-router)# bgp log-neighbor-changes`	(Optional) Enables the generation of logging messages when the status of a BGP neighbor changes. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 5	**bgp update-delay** *time-period* **Example:** `Device(config-router)# bgp update-delay 1`	(Optional) Sets the maximum initial delay period before sending the first update. The range is 1 to 3600 seconds. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 6	**bgp graceful-restart** **Example:** `Device(config-router)# bgp graceful-restart`	(Optional) Enables the BGP graceful restart capability for all BGP neighbors. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 7	**no bgp default ipv4-unicast** **Example:** `Device(config-router)# no bgp default ipv4-unicast`	(Optional) Disables default IPv4 unicast address family for BGP peering session establishment. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 8	**neighbor** *ip-address* **remote-as** *number* **Example:** `Device(config-router)# neighbor 10.11.11.11 remote-as 1`	Defines multiprotocol-BGP neighbors. Under each neighbor, define the Layer 2 Virtual Private Network (L2VPN) EVPN configuration. Use the IP address of the spine switch as the neighbor IP address.

	Command or Action	Purpose
Step 9	**neighbor** {*ip-address* | *group-name*} **update-source** *interface* **Example:** `Device(config-router)# neighbor 10.11.11.11 update-source Loopback0`	Configures update source. Update source can be configured per neighbor or per peer-group. Use the IP address of the spine switch as the neighbor IP address.
Step 10	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Specifies the L2VPN address family and enters address family configuration mode.
Step 11	**neighbor** *ip-address* **activate** **Example:** `Device(config-router-af)# neighbor 10.11.11.11 activate`	Enables the exchange information from a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 12	**neighbor** *ip-address* **send-community** [**both** | **extended** | **standard**] **Example:** `Device(config-router-af)# neighbor 10.11.11.11 send-community both`	Specifies the communities attribute sent to a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 13	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 14	**end** **Example:** `Device(config-router)# end`	Returns to privileged EXEC mode.

Verifying EVPN VXLAN Layer 2 Overlay Network

The following table lists the **show** commands that are used to verify a Layer 2 VXLAN overlay network:

Table 1: Commands to Verify EVPN VXLAN Layer 2 Overlay Network

Command	Purpose
show l2vpn evpn evi [**detail**]	Displays detailed information for a particular EVPN instance or all EVPN instances.
show l2vpn evpn mac [**detail**]	Displays the MAC address database for Layer 2 EVPN.
show l2vpn evpn mac ip [**detail**]	Displays the IP address database for Layer 2 EVPN.

Command	Purpose
show l2vpn evpn summary	Displays a summary of Layer 2 EVPN information.
show l2vpn evpn capabilities	Displays platform capability information for Layer 2 EVPN.
show l2vpn evpn peers	Displays Layer 2 EVPN peer route counts and up time.
show l2vpn evpn route-target	Displays Layer 2 EVPN import route targets.
show l2vpn evpn memory	Displays Layer 2 EVPN memory usage.
show l2route evpn summary	Displays a summary of EVPN routes.
show l2route evpn mac [**detail**]	Displays MAC address information learnt by the switch in the EVPN control plane.
show l2route evpn mac ip [**detail**]	Displays MAC and IP address information learnt by the switch in the EVPN control plane.
show l2route evpn imet detail	Displays the IMET route details for Layer 2 EVPN address family. This command shows details only about traffic forwarded using ingress replication.
show bgp l2vpn evpn	Displays BGP information for Layer 2 VPN EVPN address family.
show bgp l2vpn evpn route-type 2	Displays BGP information for route type 2 of L2VPN EVPN address family.
show bgp l2vpn evpn evi context	Displays context information for Layer 2 EVPN instances.
show bgp l2vpn evpn evi *evpn-instance-id* **route-type 3**	Displays route type 3 information for the specified Layer 2 EVPN instance. This command shows details only about traffic forwarded using ingress replication.
show l2fib bridge-domain *bridge-domain-number* **detail**	Displays detailed information for a Layer 2 forwarding information base bridge domain.
show l2fib bridge-domain *bridge-domain-number* **address unicast**	Displays unicast MAC address information for a Layer 2 forwarding information base bridge domain.
show nve vni	Displays information about VXLAN network identifier members associated with an NVE interface.
show nve vni *vni-id* **detail**	Displays detailed NVE interface state information for a VXLAN network identifier member.

Command	Purpose
show nve peers	Displays NVE interface state information for peer leaf switches.
show mac address-table vlan *vlan-id*	Displays MAC addresses for a VLAN.
show platform software fed switch active matm macTable vlan *vlan-id*	Displays MAC addresses for a VLAN from MAC address table manager database for Forwarding Engine Driver (FED).
show device-tracking database	Displays device tracking database.
show device-tracking database mac	Displays device tracking MAC address database.
show ip mroute	Displays multicast routing table information.
show ip bgp l2vpn evpn detail *l2vpn-evpn-route*	Displays detailed information about a specific route.
show ip bgp l2vpn evpn detail [*mac-address* | *ip-address*]	Display routes containing an IP address only or both MAC address and IP address.
show ip bgp l2vpn evpn route-type 2 *ethernet-tag* { *mac-address* }	Displays other MAC address formats for route-type 2 EVPN routes.

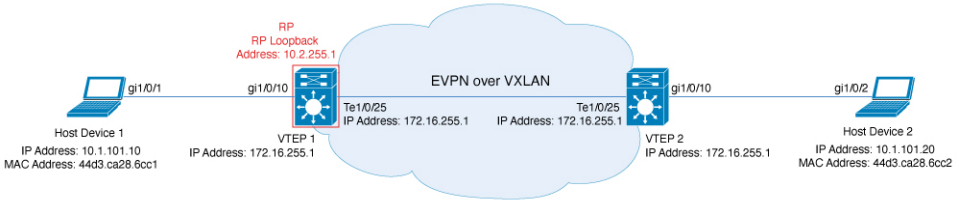
Configuration Examples for EVPN VXLAN Layer 2 Overlay Network

This sections provides configuration examples for EVPN VXLAN Layer 2 Overlay Network:

Example: Configuring Layer 2 VNI with Back-to-Back Multicast Replication

This example shows how to configure and verify a Layer 2 VNI with back-to-back multicast replication using the following topology:

Figure 7: EVPN VXLAN Network with a Layer 2 VNI with Multicast Replication



The topology shows an EVPN VXLAN network with two VTEPs (VTEP 1 and VTEP 2) and no spine switches. Multicast replication is performed between the VTEPs to forward BUM traffic in the network. VTEP 1 acts as the rendezvous point (RP) for the multicast BUM traffic. The following table provides sample configurations for the devices in this topology:

Note In a two-VTEP topology, a spine switch is not mandatory. For information about configuration of spine switches in an EVPN VXLAN network, see *Configuring Spine Switches in a BGP EVPN VXLAN Fabric* module.

Table 2: Configuring VTEP 1 and VTEP 2 to Configure a Layer 2 VNI with Back-to-Back Multicast Replication

VTEP 1	VTEP 2
Device# **show running-config** hostname Device ! ip routing ! ip multicast-routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Gi1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface Te1/0/25 no switchport ip address 172.16.12.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 !	Device# **show running-config** hostname Device ! ip routing ! ip multicast-routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.2 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Gi1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface Te1/0/25 no switchport ip address 172.16.12.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 !

VTEP 1	VTEP 2
router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-address 172.16.255.1 ! end Device#	router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both exit-address-family ! ip pim rp-address 172.16.255.1 ! end Device#

Verifying the Layer 2 VNI with Back-to-Back Multicast Replication

The following sections provide sample outputs for **show** commands to verify the Layer 2 VNI with back-to-back multicast replication on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 29
- Outputs to Verify the Configuration on VTEP 2, on page 32

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.2     2              10101      UP   N/A  00:37:39

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 1:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 7, main routing table version 7
6 network entries using 2304 bytes of memory
6 path entries using 1272 bytes of memory
2/2 BGP path/bestpath attribute entries using 576 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4192 total bytes of memory
BGP activity 6/0 prefixes, 6/0 paths, scan interval 60 secs
6 networks peaked at 10:04:33 Oct 26 2020 UTC (00:37:39.064 ago)
```

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001      45      47        7    0    0 00:38:49        2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 1:

```
Device# show bgp l2vpn evpn
BGP table version is 7, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.1:101
 *>   [2][172.16.254.1:101][0][48][44D3CA286CC1][0][*]/20
                       ::                                    32768 ?
 *>   [2][172.16.254.1:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                       ::                                    32768 ?
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.2             0    100      0 ?
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                       172.16.254.2             0    100      0 ?
Route Distinguisher: 172.16.254.2:101
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.2             0    100      0 ?
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                       172.16.254.2             0    100      0 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 1:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          Gi1/0/10:101
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          172.16.254.2

Device#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Device# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
```

```
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 00:46:14/00:03:14, RP 172.16.255.1, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Te1/0/25, Forward/Sparse, 00:43:31/00:03:14
    Loopback0, Forward/Sparse, 00:46:14/00:02:42

(*, 225.0.0.101), 00:46:14/stopped, RP 172.16.255.1, flags: SJCFx
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Te1/0/25, Forward/Sparse, 00:43:31/00:03:17
    Tunnel0, Forward/Sparse-Dense, 00:46:14/00:01:47

(172.16.254.1, 225.0.0.101), 00:00:00/00:02:59, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    Te1/0/25, Forward/Sparse, 00:00:00/00:03:29

(172.16.254.2, 225.0.0.101), 00:00:03/00:02:56, flags: x
  Incoming interface: Te1/0/25, RPF nbr 172.16.12.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:03/00:02:56

Device#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Device# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel2 Flags: A
   Te1/0/25 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
```

```
   SW Forwarding: 2/0/96/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
   Te1/0/25 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (172.16.254.1,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/96/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Null0 Flags: A
   Te1/0/25 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.2,225.0.0.101) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Te1/0/25 Flags: NS

Device#
```

Return to Verifying the Layer 2 VNI with Back-to-Back Multicast Replication, on page 29.

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.1     2              10101      UP   N/A  00:38:32

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 2:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 7, main routing table version 7
6 network entries using 2304 bytes of memory
6 path entries using 1272 bytes of memory
2/2 BGP path/bestpath attribute entries using 576 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4192 total bytes of memory
BGP activity 6/0 prefixes, 6/0 paths, scan interval 60 secs
6 networks peaked at 10:02:19 Oct 26 2020 UTC (00:38:32.591 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001      48      46        7    0    0 00:39:42        2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 2:

```
Device# show bgp l2vpn evpn
BGP table version is 7, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
```

```
             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
             x best-external, a additional-path, c RIB-compressed,
             t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.1:101
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.1             0    100      0 ?
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.1             0    100      0 ?
Route Distinguisher: 172.16.254.2:101
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.1             0    100      0 ?
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.1             0    100      0 ?
 *>   [2][172.16.254.2:101][0][48][44D3CA286CC2][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.2:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      ::                                    32768 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 2:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          172.16.254.1
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          Gi1/0/10:101

Device#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Device# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 00:43:49/00:02:09, RP 172.16.255.1, flags: SJCL
  Incoming interface: Te1/0/25, RPF nbr 172.16.12.1
  Outgoing interface list:
    Loopback0, Forward/Sparse, 00:43:49/00:02:09
```

```
(*, 225.0.0.101), 00:43:49/stopped, RP 172.16.255.1, flags: SJCFx
  Incoming interface: Te1/0/25, RPF nbr 172.16.12.1
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:43:49/00:01:11

(172.16.254.1, 225.0.0.101), 00:00:17/00:02:42, flags: JTx
  Incoming interface: Te1/0/25, RPF nbr 172.16.12.1
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:17/00:02:42

(172.16.254.2, 225.0.0.101), 00:00:20/00:02:39, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0, Registering
  Outgoing interface list:
    Te1/0/25, Forward/Sparse, 00:00:20/00:03:09

Device#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Device# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Te1/0/25 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/141/0, Other: 0/0/0
   Te1/0/25 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.1,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/96/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Te1/0/25 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.2,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/96/0, Other: 0/0/0
   HW Forwarding:   1/0/114/0, Other: 0/0/0
   Null0 Flags: A
   Te1/0/25 Flags: F NS
```

```
    Pkts: 0/0/0    Rate: 0 pps
  Tunnel1 Flags: F
    Pkts: 0/0/1    Rate: 0 pps

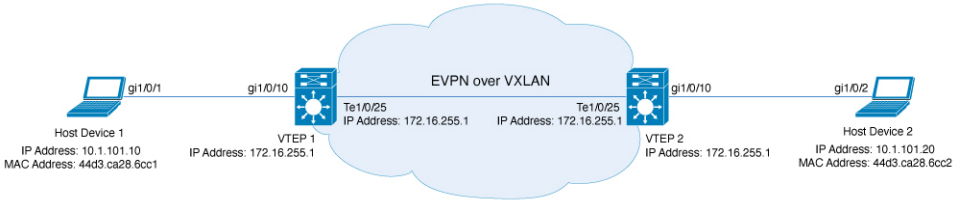
Device#
```

Return to Verifying the Layer 2 VNI with Back-to-Back Multicast Replication, on page 29.

Example: Configuring Layer 2 VNI with Back to Back Ingress Replication

This example shows how to configure and verify a Layer 2 VNI with back-to-back ingress replication using the following topology:

Figure 8: EVPN VXLAN Network with a Layer 2 VNI with Ingress Replication



The topology shows an EVPN VXLAN network with two VTEPs (VTEP 1 and VTEP 2) and no spine switches. Ingress replication is performed between the VTEPs to forward BUM traffic in the network. The following table provides sample configurations for the devices in this topology:

Note In a two-VTEP topology, a spine switch is not mandatory. For information about configuration of spine switches in an EVPN VXLAN network, see *Configuring Spine Switches in a BGP EVPN VXLAN Fabric* module.

Table 3: Configuring VTEP 1 and VTEP 2 to Configure a Layer 2 VNI with Back-to-Back Ingress Replication

VTEP 1	VTEP 2
Device# **show running-config** hostname Device ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.1 255.255.255.255 ip ospf 1 area 0 ! interface gi1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface Te1/0/25 no switchport ip address 172.16.12.1 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication !	Device# **show running-config** hostname Device ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.2 255.255.255.255 ip ospf 1 area 0 ! interface gi1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface Te1/0/25 no switchport ip address 172.16.12.2 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication !

VTEP 1	VTEP 2
router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end Device#	router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both exit-address-family ! end Device#

Verifying the Layer 2 VNI with Back-to-Back Ingress Replication

The following sections provide sample outputs for **show** commands to verify the Layer 2 VNI with back-to-back ingress replication on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 29
- Outputs to Verify the Configuration on VTEP 2, on page 32

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.2     3              10101      UP   N/A  00:34:36

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 1:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 34, main routing table version 34
9 network entries using 3456 bytes of memory
9 path entries using 1908 bytes of memory
4/4 BGP path/bestpath attribute entries using 1152 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6556 total bytes of memory
BGP activity 13/4 prefixes, 23/14 paths, scan interval 60 secs
9 networks peaked at 12:35:03 Oct 26 2020 UTC (00:34:37.010 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
```

```
172.16.255.2    4        65001     213     215       34    0    0 03:06:17        3

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 1:

```
Device# show bgp l2vpn evpn
BGP table version is 34, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.1:101
 *>   [2][172.16.254.1:101][0][48][44D3CA286CC1][0][*]/20
                       ::                                    32768 ?
 *>   [2][172.16.254.1:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                       ::                                    32768 ?
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.2             0    100      0 ?
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                       172.16.254.2             0    100      0 ?
Route Distinguisher: 172.16.254.2:101
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.2             0    100      0 ?
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                       172.16.254.2             0    100      0 ?
Route Distinguisher: 172.16.254.1:101
 *>   [3][172.16.254.1:101][0][32][172.16.254.1]/17
                       ::                                    32768 ?
 *>i  [3][172.16.254.1:101][0][32][172.16.254.2]/17
                       172.16.254.2             0    100      0 ?
Route Distinguisher: 172.16.254.2:101
 *>i  [3][172.16.254.2:101][0][32][172.16.254.2]/17
                       172.16.254.2             0    100      0 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 1:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ---------------------------- ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          Gi1/0/10:101
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          172.16.254.2

Device#
```

The following example shows the output for the **show l2fib bridge-domain** *evpn-instance* **detail** command on VTEP 1:

```
Device# show l2fib bridge-domain 101 detail
Bridge Domain : 101
  Reference Count : 10
  Replication ports count : 2
  Unicast Address table size : 1
  IP Multicast Prefix table size : 3
```

```
  Flood List Information :
    Olist: 1125, Ports: 2

  Port Information :
    BD_PORT   Gi1/0/10:101
    VXLAN_REP PL:25(1) T:VXLAN_REP [IR]10101:172.16.254.2

  Unicast Address table information :
    44d3.ca28.6cc2  VXLAN_UC  PL:24(1) T:VXLAN_UC [MAC]10101:172.16.254.2

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 1125, Ports: 2

Device#
```

Return to Verifying the Layer 2 VNI with Back-to-Back Multicast Replication, on page 29.

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.1     3              10101      UP   N/A  00:35:22

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 2:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 34, main routing table version 34
9 network entries using 3456 bytes of memory
9 path entries using 1908 bytes of memory
4/4 BGP path/bestpath attribute entries using 1152 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6556 total bytes of memory
BGP activity 13/4 prefixes, 23/14 paths, scan interval 60 secs
9 networks peaked at 12:32:49 Oct 26 2020 UTC (00:34:55.476 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     215     213       34    0    0 03:06:35        3

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 2:

```
Device# show bgp l2vpn evpn
BGP table version is 34, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
```

```
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.1:101
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.1             0    100      0 ?
 *>i  [2][172.16.254.1:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.1             0    100      0 ?
Route Distinguisher: 172.16.254.2:101
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.1             0    100      0 ?
 *>i  [2][172.16.254.2:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.1             0    100      0 ?
 *>   [2][172.16.254.2:101][0][48][44D3CA286CC2][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.2:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      ::                                    32768 ?
Route Distinguisher: 172.16.254.1:101
 *>i  [3][172.16.254.1:101][0][32][172.16.254.1]/17
                      172.16.254.1             0    100      0 ?
Route Distinguisher: 172.16.254.2:101
 *>i  [3][172.16.254.2:101][0][32][172.16.254.1]/17
                      172.16.254.1             0    100      0 ?
 *>   [3][172.16.254.2:101][0][32][172.16.254.2]/17
                      ::                                    32768 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 2:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          172.16.254.1
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          gi1/0/10

Device#
```

The following example shows the output for the **show l2fib bridge-domain** *evpn-instance* **detail** command on VTEP 2:

```
Device# show l2fib bridge-domain 101 detail
Bridge Domain : 101
  Reference Count : 10
  Replication ports count : 2
  Unicast Address table size : 1
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 1125, Ports: 2

  Port Information :
    BD_PORT   Gi1/0/10:101
    VXLAN_REP PL:16(1) T:VXLAN_REP [IR]10101:172.16.254.1

  Unicast Address table information :
    44d3.ca28.6cc1  VXLAN_UC  PL:15(1) T:VXLAN_UC [MAC]10101:172.16.254.1

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 1125, Ports: 2
```

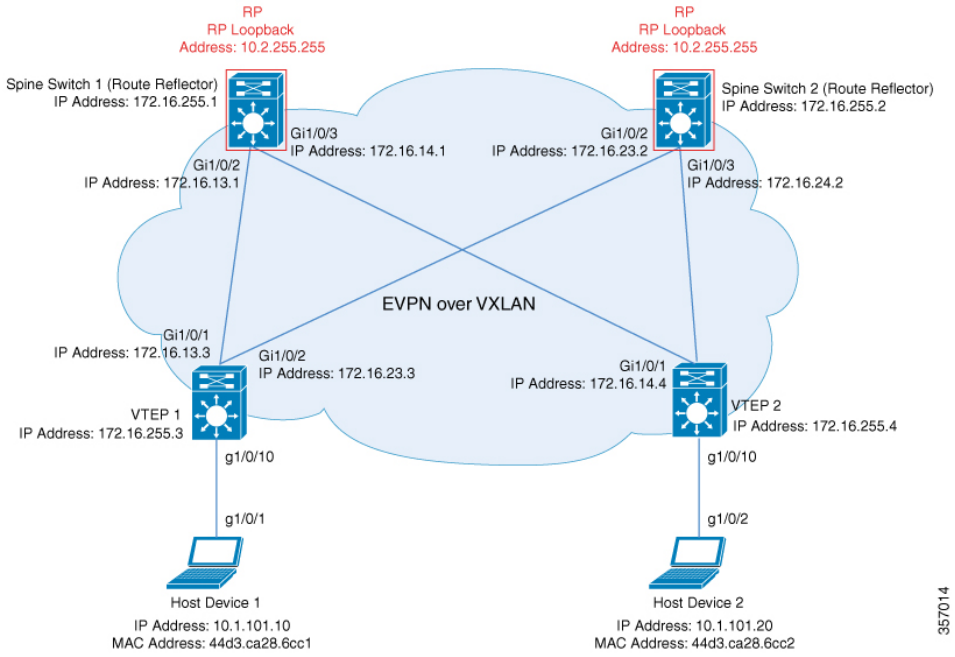
```
Device#
```

Return to Verifying the Layer 2 VNI with Back-to-Back Multicast Replication, on page 29.

Example: Configuring Layer 2 VNI with Spine Multicast Replication

This example shows how to configure and verify a Layer 2 VNI with spine multicast replication using the following topology:

Figure 9: EVPN VXLAN Network with a Layer 2 VNI with Multicast Replication



The topology shows an EVPN VXLAN network with two spine switches (Spine Switch 1 and Spine Switch 2) and two VTEPs (VTEP 1 and VTEP 2). Multicast replication is performed between the VTEPs to forward BUM traffic in the network. Spine Switch 1 and Spine Switch 2 act as route reflectors and also as the RPs for the multicast BUM traffic in the network. The following tables provide sample configurations for the devices in this topology:

Table 4: Configuring VTEP 1 and VTEP 2 to Configure a Layer 2 VNI with Spine Multicast Replication

VTEP 1	VTEP 2
Device# **show running-config** hostname Device ! ip routing ! ip multicast-routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 ! router ospf 1 router-id 172.16.255.3 !	Device# **show running-config** hostname Device ! ip routing ! ip multicast-routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 ! router ospf 1 router-id 172.16.255.4 !

VTEP 1	VTEP 2
router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-address 172.16.255.255 ! end Device#	router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-address 172.16.255.255 ! end Device#

Table 5: Configuring VTEP 1 and VTEP 2 to Configure a Layer 2 VNI with Spine Multicast Replication

Spine Switch 1	Spine Switch 2
Device# **show running-config** hostname Device ! ip routing ! ip multicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 no switchport ip address 172.16.14.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family !	Device# **show running-config** hostname Device ! ip routing ! ip multicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.2 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 no switchport ip address 172.16.24.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp router-id 172.16.255.2 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family !

Spine Switch 1	Spine Switch 2
address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end Device#	address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end Device#

Verifying the Layer 2 VNI with Spine Multicast Replication

The following sections provide sample outputs for **show** commands to verify the Layer 2 VNI with spine multicast replication on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 45
- Outputs to Verify the Configuration on VTEP 2, on page 49
- Outputs to Verify the Configuration on Spine Switch 1 (RP inside the Network), on page 52
- Outputs to Verify the Configuration on Spine Switch 2 (RP inside the Network), on page 56

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.4     2              10101      UP   N/A  00:00:56

Device#
```

The following example shows the output for the **show ip route** command on VTEP 1:

```
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected
```

```
Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 15 subnets, 2 masks
C        172.16.13.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.13.3/32 is directly connected, GigabitEthernet1/0/1
O        172.16.14.0/24
           [110/2] via 172.16.13.1, 01:43:35, GigabitEthernet1/0/1
C        172.16.23.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.23.3/32 is directly connected, GigabitEthernet1/0/2
O        172.16.24.0/24
           [110/2] via 172.16.23.2, 01:43:35, GigabitEthernet1/0/2
O        172.16.254.1/32
           [110/2] via 172.16.13.1, 00:09:33, GigabitEthernet1/0/1
O        172.16.254.2/32
           [110/2] via 172.16.23.2, 00:08:17, GigabitEthernet1/0/2
C        172.16.254.3/32 is directly connected, Loopback1
O        172.16.254.4/32
           [110/3] via 172.16.23.2, 01:43:35, GigabitEthernet1/0/2
           [110/3] via 172.16.13.1, 01:43:35, GigabitEthernet1/0/1
O        172.16.255.1/32
           [110/2] via 172.16.13.1, 01:43:35, GigabitEthernet1/0/1
O        172.16.255.2/32
           [110/2] via 172.16.23.2, 01:43:35, GigabitEthernet1/0/2
C        172.16.255.3/32 is directly connected, Loopback0
O        172.16.255.4/32
           [110/3] via 172.16.23.2, 01:43:35, GigabitEthernet1/0/2
           [110/3] via 172.16.13.1, 01:43:35, GigabitEthernet1/0/1
O        172.16.255.255/32
           [110/2] via 172.16.23.2, 00:08:17, GigabitEthernet1/0/2
           [110/2] via 172.16.13.1, 00:09:33, GigabitEthernet1/0/1

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 1:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 54, main routing table version 54
6 network entries using 2304 bytes of memory
8 path entries using 1696 bytes of memory
2/2 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4696 total bytes of memory
BGP activity 15/9 prefixes, 33/25 paths, scan interval 60 secs
9 networks peaked at 16:10:51 Oct 26 2020 UTC (01:42:36.958 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     133     120       54    0    0 01:43:34        2
172.16.255.2    4        65001     134     123       54    0    0 01:43:34        2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 1:

```
Device# show bgp l2vpn evpn
BGP table version is 54, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>   [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      ::                                    32768 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 1:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          Gi1/0/
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          172.16.254.4

Device#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Device# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 00:05:22/00:02:42, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
```

```
    Loopback1, Forward/Sparse, 00:05:20/00:02:42

(*, 225.0.0.101), 00:01:34/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:01:34/00:01:27

(172.16.254.4, 225.0.0.101), 00:00:57/00:02:02, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:57/00:02:02

(172.16.254.3, 225.0.0.101), 00:01:32/00:01:27, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0, Registering
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:32/00:02:57

Device#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Device# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback1 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/114/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 1/1/0
   HW Forwarding:   148/0/155/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Tunnel1 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.101) Flags: HW
```

```
   SW Forwarding: 1/0/96/0, Other: 0/0/0
   HW Forwarding:   2/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps

Device#
```

Return to Verifying the Layer 2 VNI with Spine Multicast Replication, on page 45.

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.3     2              10101      UP   N/A  00:01:39

Device#
```

The following example shows the output for the **show ip route** command on VTEP 2:

```
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 15 subnets, 2 masks
O        172.16.13.0/24
           [110/2] via 172.16.14.1, 01:44:23, GigabitEthernet1/0/1
C        172.16.14.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.14.4/32 is directly connected, GigabitEthernet1/0/1
O        172.16.23.0/24
           [110/2] via 172.16.24.2, 01:44:23, GigabitEthernet1/0/2
C        172.16.24.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.24.4/32 is directly connected, GigabitEthernet1/0/2
O        172.16.254.1/32
           [110/2] via 172.16.14.1, 00:10:18, GigabitEthernet1/0/1
O        172.16.254.2/32
           [110/2] via 172.16.24.2, 00:09:02, GigabitEthernet1/0/2
O        172.16.254.3/32
           [110/3] via 172.16.24.2, 01:44:20, GigabitEthernet1/0/2
           [110/3] via 172.16.14.1, 01:44:15, GigabitEthernet1/0/1
C        172.16.254.4/32 is directly connected, Loopback1
O        172.16.255.1/32
           [110/2] via 172.16.14.1, 01:44:23, GigabitEthernet1/0/1
O        172.16.255.2/32
           [110/2] via 172.16.24.2, 01:44:23, GigabitEthernet1/0/2
O        172.16.255.3/32
           [110/3] via 172.16.24.2, 01:44:20, GigabitEthernet1/0/2
           [110/3] via 172.16.14.1, 01:44:15, GigabitEthernet1/0/1
```

```
C        172.16.255.4/32 is directly connected, Loopback0
O        172.16.255.255/32
           [110/2] via 172.16.24.2, 00:09:01, GigabitEthernet1/0/2
           [110/2] via 172.16.14.1, 00:10:18, GigabitEthernet1/0/1

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 2:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 54, main routing table version 54
6 network entries using 2304 bytes of memory
8 path entries using 1696 bytes of memory
2/2 BGP path/bestpath attribute entries using 576 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4696 total bytes of memory
BGP activity 15/9 prefixes, 34/26 paths, scan interval 60 secs
9 networks peaked at 16:08:37 Oct 26 2020 UTC (01:43:22.226 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     134     123       54    0    0 01:44:22        2
172.16.255.2    4        65001     134     123       54    0    0 01:44:15        2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 2:

```
Device# show bgp l2vpn evpn
BGP table version is 54, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 *>   [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      ::                                    32768 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 2:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          Gi1/0/10:101

Device#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Device# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 00:05:51/00:02:24, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Loopback1, Forward/Sparse, 00:05:49/00:02:09
    GigabitEthernet1/0/2, Forward/Sparse, 00:05:43/00:02:24

(*, 225.0.0.101), 00:02:46/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:02:46/00:00:15

(172.16.254.4, 225.0.0.101), 00:01:43/00:01:16, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:43/00:02:45

(172.16.254.3, 225.0.0.101), 00:02:19/00:00:40, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:02:19/00:00:40

Device#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Device# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
```

```
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A NS
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback1 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   146/0/167/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.4,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/96/0, Other: 1/1/0
   HW Forwarding:   4/0/145/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps

Device#
```

Return to Verifying the Layer 2 VNI with Spine Multicast Replication, on page 45.

Outputs to Verify the Configuration on Spine Switch 1 (RP inside the Network)

The following example shows the output for the **show ip route** command on Spine Switch 1:

```
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
```

```
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 15 subnets, 2 masks
C        172.16.13.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.13.1/32 is directly connected, GigabitEthernet1/0/2
C        172.16.14.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.14.1/32 is directly connected, GigabitEthernet1/0/3
O        172.16.23.0/24
           [110/2] via 172.16.13.3, 01:45:08, GigabitEthernet1/0/2
O        172.16.24.0/24
           [110/2] via 172.16.14.4, 01:45:12, GigabitEthernet1/0/3
C        172.16.254.1/32 is directly connected, Loopback1
O        172.16.254.2/32
           [110/3] via 172.16.14.4, 00:09:51, GigabitEthernet1/0/3
           [110/3] via 172.16.13.3, 00:09:51, GigabitEthernet1/0/2
O        172.16.254.3/32
           [110/2] via 172.16.13.3, 01:45:08, GigabitEthernet1/0/2
O        172.16.254.4/32
           [110/2] via 172.16.14.4, 01:45:12, GigabitEthernet1/0/3
C        172.16.255.1/32 is directly connected, Loopback0
O        172.16.255.2/32
           [110/3] via 172.16.14.4, 01:45:12, GigabitEthernet1/0/3
           [110/3] via 172.16.13.3, 01:45:08, GigabitEthernet1/0/2
O        172.16.255.3/32
           [110/2] via 172.16.13.3, 01:45:08, GigabitEthernet1/0/2
O        172.16.255.4/32
           [110/2] via 172.16.14.4, 01:45:12, GigabitEthernet1/0/3
C        172.16.255.255/32 is directly connected, Loopback2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 1:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 35, main routing table version 35
4 network entries using 1376 bytes of memory
8 path entries using 1664 bytes of memory
1/1 BGP path/bestpath attribute entries using 288 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3448 total bytes of memory
BGP activity 12/8 prefixes, 28/20 paths, scan interval 60 secs
6 networks peaked at 16:08:39 Oct 26 2020 UTC (01:44:10.445 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001     133     132       35    0    0 01:45:07        4
172.16.255.3    4        65001     122     135       35    0    0 01:45:07        2
172.16.255.4    4        65001     124     135       35    0    0 01:45:10        2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 1:

```
Device# show bgp l2vpn evpn
BGP table version is 35, local router ID is 172.16.255.1
```

```
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?

Device#
```

The following example shows the output for the **show ip msdp summary** command on Spine Switch 1:

```
Device# show ip msdp summary
MSDP Peer Status Summary
Peer Address     AS    State    Uptime/  Reset SA    Peer Name
                                Downtime Count Count
172.16.254.2     65001 Up       00:06:28 0     0     ?

Device#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 1:

```
Device# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 00:56:14/00:02:21, RP 172.16.255.255, flags: SPL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(*, 225.0.0.101), 00:00:12/stopped, RP 172.16.255.255, flags: SP
```

```
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.4, 225.0.0.101), 00:00:05/00:02:54, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.14.4
  Outgoing interface list: Null

(172.16.254.3, 225.0.0.101), 00:00:12/00:02:47, flags: PA
  Incoming interface: GigabitEthernet1/0/3, RPF nbr 172.16.13.3
  Outgoing interface list: Null

Device#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 1:

```
Device# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 2/2/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: IC
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 1/0/1
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: NS
 (172.16.254.4,225.0.0.101) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/3 Flags: NS

Device#
```

Return to Verifying the Layer 2 VNI with Spine Multicast Replication, on page 45.

Outputs to Verify the Configuration on Spine Switch 2 (RP inside the Network)

The following example shows the output for the **show ip route** command on Spine Switch 2:

```
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 15 subnets, 2 masks
O        172.16.13.0/24
           [110/2] via 172.16.23.3, 01:45:34, GigabitEthernet1/0/2
O        172.16.14.0/24
           [110/2] via 172.16.24.4, 01:45:38, GigabitEthernet1/0/3
C        172.16.23.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.23.2/32 is directly connected, GigabitEthernet1/0/2
C        172.16.24.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.24.2/32 is directly connected, GigabitEthernet1/0/3
O        172.16.254.1/32
           [110/3] via 172.16.24.4, 00:11:33, GigabitEthernet1/0/3
           [110/3] via 172.16.23.3, 00:11:33, GigabitEthernet1/0/2
C        172.16.254.2/32 is directly connected, Loopback1
O        172.16.254.3/32
           [110/2] via 172.16.23.3, 01:45:34, GigabitEthernet1/0/2
O        172.16.254.4/32
           [110/2] via 172.16.24.4, 01:45:38, GigabitEthernet1/0/3
O        172.16.255.1/32
           [110/3] via 172.16.24.4, 01:45:34, GigabitEthernet1/0/3
           [110/3] via 172.16.23.3, 01:45:30, GigabitEthernet1/0/2
C        172.16.255.2/32 is directly connected, Loopback0
O        172.16.255.3/32
           [110/2] via 172.16.23.3, 01:45:34, GigabitEthernet1/0/2
O        172.16.255.4/32
           [110/2] via 172.16.24.4, 01:45:38, GigabitEthernet1/0/3
C        172.16.255.255/32 is directly connected, Loopback2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 2:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 35, main routing table version 35
4 network entries using 1376 bytes of memory
8 path entries using 1664 bytes of memory
1/1 BGP path/bestpath attribute entries using 288 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3448 total bytes of memory
BGP activity 10/6 prefixes, 28/20 paths, scan interval 60 secs
```

```
6 networks peaked at 16:09:46 Oct 26 2020 UTC (01:44:35.591 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     133     134       35    0    0 01:45:33        4
172.16.255.3    4        65001     125     137       35    0    0 01:45:33        2
172.16.255.4    4        65001     125     136       35    0    0 01:45:28        2

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 2:

```
Device# show bgp l2vpn evpn
BGP table version is 35, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?

Device#
```

The following example shows the output for the **show ip msdp summary** command on Spine Switch 2:

```
Device# show ip msdp summary
MSDP Peer Status Summary
Peer Address     AS    State    Uptime/  Reset SA    Peer Name
                                Downtime Count Count
172.16.254.1     65001 Up       00:06:53 0     2     ?

Device#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 2:

```
Device# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
```

```
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 00:56:18/00:03:26, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:54:14/00:03:08
    GigabitEthernet1/0/3, Forward/Sparse, 00:56:18/00:03:26

(*, 225.0.0.101), 00:51:00/00:03:17, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:50:34/00:03:17
    GigabitEthernet1/0/3, Forward/Sparse, 00:51:00/00:02:43

(172.16.254.4, 225.0.0.101), 00:00:17/00:02:42, flags: TA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/3, Forward/Sparse, 00:00:17/00:03:17

(172.16.254.3, 225.0.0.101), 00:00:23/00:02:36, flags: TA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/3, Forward/Sparse, 00:00:23/00:03:06

Device#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 2:

```
Device# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/3 Flags: F NS
```

```
    Pkts: 0/0/0    Rate: 0 pps
(*,225.0.0.101) Flags: C HW
  SW Forwarding: 2/0/150/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
  Tunnel1 Flags: A
  GigabitEthernet1/0/2 Flags: F NS
    Pkts: 0/0/2    Rate: 0 pps
  GigabitEthernet1/0/3 Flags: F NS
    Pkts: 0/0/2    Rate: 0 pps
(172.16.254.3,225.0.0.101) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A
  GigabitEthernet1/0/3 Flags: F NS
    Pkts: 0/0/0    Rate: 0 pps
(172.16.254.4,225.0.0.101) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A
  GigabitEthernet1/0/3 Flags: F NS
    Pkts: 0/0/0    Rate: 0 pps

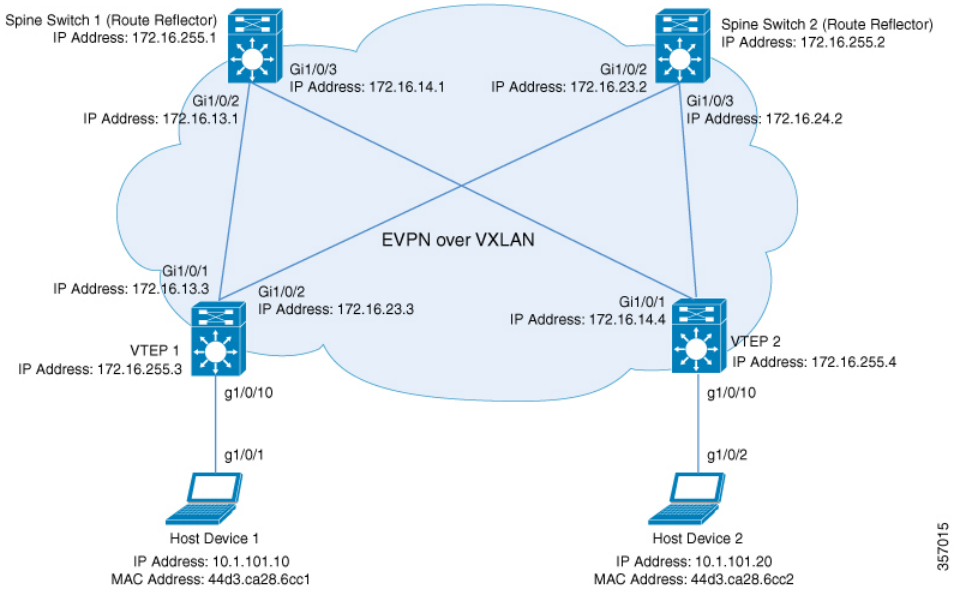
Device#
```

Return to Verifying the Layer 2 VNI with Spine Multicast Replication, on page 45.

Example: Configuring Layer 2 VNI with Spine Ingress Replication

This example shows how to configure and verify a Layer 2 VNI with spine ingress replication using the following topology:

Figure 10: EVPN VXLAN Network with a Layer 2 VNI with Ingress Replication



The topology shows an EVPN VXLAN network with two spine switches (Spine Switch 1 and Spine Switch 2) and two VTEPs (VTEP 1 and VTEP 2). Ingress replication is performed between the VTEPs to forward

BUM traffic in the network. Spine Switch 1 and Spine Switch 2 act as route reflectors in the network. The following tables provide sample configurations for the devices in this topology:

Table 6: Configuring VTEP 1 and VTEP 2 to Configure a Layer 2 VNI with Spine Ingress Replication

VTEP 1	VTEP 2
Device# **show running-config** hostname Device ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication ! router ospf 1 router-id 172.16.255.3 !	Device# **show running-config** hostname Device ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication ! router ospf 1 router-id 172.16.255.4 !

VTEP 1	VTEP 2
router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end Device#	router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end Device#

Table 7: Configuring VTEP 1 and VTEP 2 to Configure a Layer 2 VNI with Spine Ingress Replication

Spine Switch 1	Spine Switch 2
Spine-01# **show running-config** hostname Spine-01 ! ip routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 no switchport ip address 172.16.14.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family !	Spine-02# **show running-config** hostname Spine-02 ! ip routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 no switchport ip address 172.16.24.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp router-id 172.16.255.2 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family !
address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client exit-address-family ! end Spine-01#	address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client exit-address-family ! end Spine-02#

Verifying the Layer 2 VNI with Spine Ingress Replication

The following sections provide sample outputs for **show** commands to verify the Layer 2 VNI with spine ingress replication on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 63
- Outputs to Verify the Configuration on VTEP 2, on page 65
- Outputs to Verify the Configuration on Spine Switch 1, on page 68
- Outputs to Verify the Configuration on Spine Switch 2, on page 70

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.4     3              10101      UP   N/A  01:25:20

Device#
```

The following example shows the output for the **show ip route** command on VTEP 1:

```
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 12 subnets, 2 masks
C        172.16.13.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.13.3/32 is directly connected, GigabitEthernet1/0/1
O        172.16.14.0/24
           [110/2] via 172.16.13.1, 01:26:20, GigabitEthernet1/0/1
C        172.16.23.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.23.3/32 is directly connected, GigabitEthernet1/0/2
O        172.16.24.0/24
           [110/2] via 172.16.23.2, 01:26:20, GigabitEthernet1/0/2
C        172.16.254.3/32 is directly connected, Loopback1
O        172.16.254.4/32
           [110/3] via 172.16.23.2, 01:26:20, GigabitEthernet1/0/2
           [110/3] via 172.16.13.1, 01:26:20, GigabitEthernet1/0/1
O        172.16.255.1/32
           [110/2] via 172.16.13.1, 01:26:20, GigabitEthernet1/0/1
O        172.16.255.2/32
           [110/2] via 172.16.23.2, 01:26:20, GigabitEthernet1/0/2
C        172.16.255.3/32 is directly connected, Loopback0
O        172.16.255.4/32
           [110/3] via 172.16.23.2, 01:26:20, GigabitEthernet1/0/2
```

```
           [110/3] via 172.16.13.1, 01:26:20, GigabitEthernet1/0/1

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 1:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 13, main routing table version 13
9 network entries using 3456 bytes of memory
12 path entries using 2544 bytes of memory
4/4 BGP path/bestpath attribute entries using 1152 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7272 total bytes of memory
BGP activity 9/0 prefixes, 15/3 paths, scan interval 60 secs
9 networks peaked at 16:10:51 Oct 26 2020 UTC (01:25:22.020 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     101      99       13    0    0 01:26:19        3
172.16.255.2    4        65001     102     100       13    0    0 01:26:19        3

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 1:

```
Device# show bgp l2vpn evpn
BGP table version is 13, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>   [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                       ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                       ::                                    32768 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                       172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.4             0    100      0 ?
 *>i                    172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                       172.16.254.4             0    100      0 ?
 *>i                    172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 *>   [3][172.16.254.3:101][0][32][172.16.254.3]/17
                       ::                                    32768 ?
 *>i  [3][172.16.254.3:101][0][32][172.16.254.4]/17
                       172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
```

```
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4              0    100      0 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 1:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          Gi1/0/10:101
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          172.16.254.4

Device#
```

The following example shows the output for the **show l2fib bridge-domain** *evpn-instance* **detail** command on VTEP 1:

```
Device# show l2fib bridge-domain 101 detail
Bridge Domain : 101
  Reference Count : 10
  Replication ports count : 2
  Unicast Address table size : 1
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 1125, Ports: 2

  Port Information :
    BD_PORT   Gi1/0/10:101
    VXLAN_REP PL:2(1) T:VXLAN_REP [IR]10101:172.16.254.4

  Unicast Address table information :
    44d3.ca28.6cc2  VXLAN_UC  PL:1(1) T:VXLAN_UC [MAC]10101:172.16.254.4

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 1125, Ports: 2


Device#
```

Return to Verifying the Layer 2 VNI with Spine Ingress Replication, on page 63.

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10101    L2CP 172.16.254.3     3              10101      UP   N/A  01:27:15

Device#
```

The following example shows the output for the **show ip route** command on VTEP 2:

```
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
       & - replicated local route overrides by connected

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 12 subnets, 2 masks
O        172.16.13.0/24
           [110/2] via 172.16.14.1, 01:28:18, GigabitEthernet1/0/1
C        172.16.14.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.14.4/32 is directly connected, GigabitEthernet1/0/1
O        172.16.23.0/24
           [110/2] via 172.16.24.2, 01:28:18, GigabitEthernet1/0/2
C        172.16.24.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.24.4/32 is directly connected, GigabitEthernet1/0/2
O        172.16.254.3/32
           [110/3] via 172.16.24.2, 01:28:15, GigabitEthernet1/0/2
           [110/3] via 172.16.14.1, 01:28:10, GigabitEthernet1/0/1
C        172.16.254.4/32 is directly connected, Loopback1
O        172.16.255.1/32
           [110/2] via 172.16.14.1, 01:28:18, GigabitEthernet1/0/1
O        172.16.255.2/32
           [110/2] via 172.16.24.2, 01:28:18, GigabitEthernet1/0/2
O        172.16.255.3/32
           [110/3] via 172.16.24.2, 01:28:15, GigabitEthernet1/0/2
           [110/3] via 172.16.14.1, 01:28:10, GigabitEthernet1/0/1
C        172.16.255.4/32 is directly connected, Loopback0

Device#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 2:

```
Device# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 13, main routing table version 13
9 network entries using 3456 bytes of memory
12 path entries using 2544 bytes of memory
4/4 BGP path/bestpath attribute entries using 1152 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7272 total bytes of memory
BGP activity 9/0 prefixes, 15/3 paths, scan interval 60 secs
9 networks peaked at 16:08:37 Oct 26 2020 UTC (01:27:15.987 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     103     101       13    0    0 01:28:16        3
172.16.255.2    4        65001     103     101       13    0    0 01:28:09        3

Device#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 2:

```
Device# show bgp l2vpn evpn
BGP table version is 13, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 *>   [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      ::                                    32768 ?
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [3][172.16.254.4:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>   [3][172.16.254.4:101][0][32][172.16.254.4]/17
                      ::                                    32768 ?

Device#
```

The following example shows the output for the **show l2vpn evpn mac evi** *evpn-instance* command on VTEP 2:

```
Device# show l2vpn evpn mac evi 101
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------------ ---------- ---------------
44d3.ca28.6cc1 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
44d3.ca28.6cc2 101   101   0000.0000.0000.0000.0000 0          Gi1/0/10:101

Device#
```

The following example shows the output for the **show l2fib bridge-domain** *evpn-instance* **detail** command on VTEP 2:

```
Device# show l2fib bridge-domain 101 detail
Bridge Domain : 101
  Reference Count : 10
  Replication ports count : 2
  Unicast Address table size : 1
  IP Multicast Prefix table size : 3
```

```
  Flood List Information :
    Olist: 1125, Ports: 2

  Port Information :
    BD_PORT   Gi1/0/10:101
    VXLAN_REP PL:2(1) T:VXLAN_REP [IR]10101:172.16.254.3

  Unicast Address table information :
    44d3.ca28.6cc1  VXLAN_UC  PL:1(1) T:VXLAN_UC [MAC]10101:172.16.254.3

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 1125, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 1125, Ports: 2

Device#
```

Return to Verifying the Layer 2 VNI with Spine Ingress Replication, on page 63.

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show ip route** command on Spine Switch 1:

```
Spine-01# 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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 12 subnets, 2 masks
C        172.16.13.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.13.1/32 is directly connected, GigabitEthernet1/0/2
C        172.16.14.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.14.1/32 is directly connected, GigabitEthernet1/0/3
O        172.16.23.0/24
           [110/2] via 172.16.13.3, 01:29:42, GigabitEthernet1/0/2
O        172.16.24.0/24
           [110/2] via 172.16.14.4, 01:29:46, GigabitEthernet1/0/3
O        172.16.254.3/32
           [110/2] via 172.16.13.3, 01:29:42, GigabitEthernet1/0/2
O        172.16.254.4/32
           [110/2] via 172.16.14.4, 01:29:46, GigabitEthernet1/0/3
C        172.16.255.1/32 is directly connected, Loopback0
O        172.16.255.2/32
           [110/3] via 172.16.14.4, 01:29:46, GigabitEthernet1/0/3
           [110/3] via 172.16.13.3, 01:29:42, GigabitEthernet1/0/2
O        172.16.255.3/32
           [110/2] via 172.16.13.3, 01:29:42, GigabitEthernet1/0/2
O        172.16.255.4/32
           [110/2] via 172.16.14.4, 01:29:46, GigabitEthernet1/0/3

Spine-01#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 7, main routing table version 7
6 network entries using 2064 bytes of memory
12 path entries using 2496 bytes of memory
3/3 BGP path/bestpath attribute entries using 864 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5544 total bytes of memory
BGP activity 6/0 prefixes, 12/0 paths, scan interval 60 secs
6 networks peaked at 16:08:39 Oct 26 2020 UTC (01:28:44.518 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001     107     106        7    0    0 01:29:41        6
172.16.255.3    4        65001     102     105        7    0    0 01:29:41        3
172.16.255.4    4        65001     103     105        7    0    0 01:29:44        3

Spine-01#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn
BGP table version is 7, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?

Spine-01#
```

Return to Verifying the Layer 2 VNI with Spine Ingress Replication, on page 63.

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show ip route** command on Spine Switch 2:

```
Spine-02# 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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 12 subnets, 2 masks
O        172.16.13.0/24
           [110/2] via 172.16.23.3, 01:30:51, GigabitEthernet1/0/2
O        172.16.14.0/24
           [110/2] via 172.16.24.4, 01:30:55, GigabitEthernet1/0/3
C        172.16.23.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.23.2/32 is directly connected, GigabitEthernet1/0/2
C        172.16.24.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.24.2/32 is directly connected, GigabitEthernet1/0/3
O        172.16.254.3/32
           [110/2] via 172.16.23.3, 01:30:51, GigabitEthernet1/0/2
O        172.16.254.4/32
           [110/2] via 172.16.24.4, 01:30:55, GigabitEthernet1/0/3
O        172.16.255.1/32
           [110/3] via 172.16.24.4, 01:30:51, GigabitEthernet1/0/3
           [110/3] via 172.16.23.3, 01:30:47, GigabitEthernet1/0/2
C        172.16.255.2/32 is directly connected, Loopback0
O        172.16.255.3/32
           [110/2] via 172.16.23.3, 01:30:51, GigabitEthernet1/0/2
O        172.16.255.4/32
           [110/2] via 172.16.24.4, 01:30:55, GigabitEthernet1/0/3

Spine-02#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 7, main routing table version 7
6 network entries using 2064 bytes of memory
12 path entries using 2496 bytes of memory
3/3 BGP path/bestpath attribute entries using 864 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
1 BGP extended community entries using 40 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 5544 total bytes of memory
BGP activity 6/0 prefixes, 12/0 paths, scan interval 60 secs
6 networks peaked at 16:09:46 Oct 26 2020 UTC (01:29:52.664 ago)
```

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     108     108        7    0    0 01:30:50        6
172.16.255.3    4        65001     105     107        7    0    0 01:30:50        3
172.16.255.4    4        65001     104     106        7    0    0 01:30:46        3

Spine-02#
```

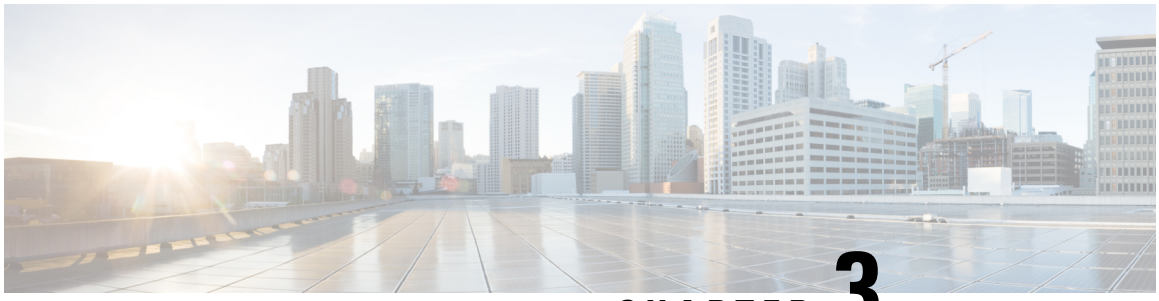
The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn
BGP table version is 7, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.10]/24
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:101][0][48][44D3CA286CC2][32][10.1.101.20]/24
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?

Spine-02#
```

Return to Verifying the Layer 2 VNI with Spine Ingress Replication, on page 63.



CHAPTER 3

BGP EVPN Layer 2 Overlay CLI Simplification

- Simplified CLIs for Layer 2 Overlay Configuration, on page 73
- Service Profile, on page 73

Simplified CLIs for Layer 2 Overlay Configuration

Use the following set of CLIs to provision an EVPN Layer 2 overlay with three contiguous VLANs and the corresponding EVPN instances (EVIs).

```
l2vpn evpn
 route-target auto vni
!
vlan configuration 100-102
 member evpn-instance
!
interface nve1
 no ip address
 source-interface Loopback2
   host-reachability protocol bgp
!
```

With the above configuration, the EVPN instance (EVI), EVI ID, and the VNI are auto-generated. The EVI ID is generated using the the EVI base and the VLAN ID. If an EVI base is not explicitly configured through a service profile, a default value of zero is assigned to it. By default, the encapsulation is set to *vxlan* and the replication-type is set to *ingress*.

Service Profile

A service profile represents a common set of characteristics that are shared by a group of services such as VLANs. A common set of characteristics that are shared by VLAN-based services can be replication-type, EVI mapping, VNI mapping, route-target mapping, and so on.

The **member evpn-instance** and **l2vpn evpn** CLIs are enhanced to provide an optional **profile** *profile-name* keyword to group the common characteristics of an EVPN instance.

l2vpn evpn profile {**default** | *profile-name*}

If you do not provide a profile name and its definition, a default EVPN profile is created. Use the **show l2vpn evpn profile default** command to see the default profile values.

```
Leaf1# show l2vpn evpn profile default detail
EVPN Profile (VLAN Based): default (auto)
  evpn id base: 0
  l2 vni base: 20000
  Encapsulation: vxlan
  Replication Type: Ingress
  IP Local Learn: Not set
  Flooding Suppression Address Resolution: Not set
  Adv. default-gateway: Not set
  Adv. Multicast: Not set
  RT5 Re-originate Flag: Disabled
  Local Routing: Not set
  vlan(s):
    101-102,111-114,131-132
Leaf1#
```

The following shows how to create an EVPN profile and use it.

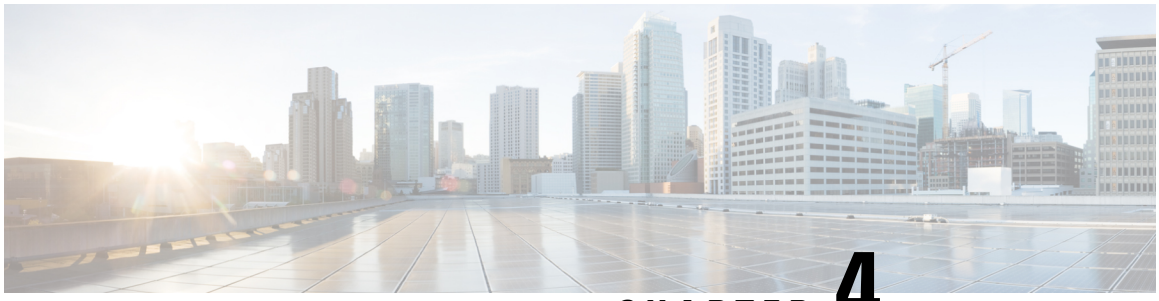
```
l2vpn evpn profile CCTV
  evi-base 3
  l2vni-base 30000
  replication-type ingress
  encapsulation vxlan
  default-gateway advertise enable
  multicast advertise enable
  ip local-learning disable }
  flooding-suppression address-resolution disable
  re-originate route-type5
  local-routing enable
!
vlan configuration 10-12
  member evpn-instance profile CCTV
```

The example creates a profile named CCTV and defines the basic EVPN characteristics like the replication type, local routing, encapsulation, flooding suppression, and so on. When this CCTV profile is applied to an EVPN instance through the **member evpn-instance profile CCTV** command, the EVPN instance inherits all the definitions under the CCTV profile.

The following table summarizes the simplified CLIs.

Replace these CLIs ….	…with New CLIs

`l2vpn evpn instance 10 vlan-based`
`  encapsulation vxlan`
`l2vpn evpn instance 11 vlan-based`
`  encapsulation vxlan`
`l2vpn evpn instance 12 vlan-based`
`  encapsulation vxlan`
`…`
`vlan configuration 10`
` member evpn-instance 10 vni 20010`
`vlan configuration 11`
` member evpn-instance 11 vni 20011`
`vlan configuration 12`
` member evpn-instance 12 vni 20012`
`…`
`interface nve1`
`  source-interface Loopback0`
`  host-reachability protocol bgp`
`member vni 20010 ingress-replication`
`  member vni 20011 ingress-replication`
`  member vni 20012 ingress-replication`
`!`
`router bgp 65001`
`…`
	`vlan configuration 10-12`
	`  member evpn-instance [profile profile-name]`
	`!`
	`interface nve1`
	` source-interface Loopback0`
	` host-reachability protocol bgp`
	`!`
	`router bgp 65001`



CHAPTER 4

Configuring EVPN VXLAN Layer 2 Overlay with Q-in-VNI

IEEE 802.1Q over Layer 2 VNI (Q-in-VNI) over EVPN VXLAN network addresses the requirement of limited network Layer 2 extension and isolation by carrying the IEEE 802.1Q tag transparently within the VXLAN header. This enables a Network-to-Network Interface (NNI) Layer 2 trunk interface with one or more 802.1Q segmented network to transport over a single Layer 2 VNI across the BGP EVPN VXLAN fabric network. Q-in-VNI provides a greater number of virtual networks to be created, with the necessary flexibility and scalability for campus network environments and other situations where a large number of Layer 2 overlays are required in a BGP EVPN VXLAN fabric.

- Restrictions for EVPN VXLAN Layer 2 Overlay with Q-in-VNI, on page 77
- Information About Q-in-VNI, on page 77
- How to Configure Q-in-VNI in a BGP EVPN VXLAN Fabric, on page 79
- Example: Configuring Q-in-VNI in a BGP EVPN VXLAN Fabric, on page 80
- Additional References for EVPN VXLAN Layer 2 Overlay with Q-in-VNI, on page 81

Restrictions for EVPN VXLAN Layer 2 Overlay with Q-in-VNI

- Selective Q-in-VNI and Q-in-VNI with VLAN mapping are not supported in this release.
- The dot1q-tunnel switchport mode access VLAN ID (S-VLAN) must be unique and should not be assigned to trunk VLAN ID (C-VLAN).
- Layer 2 Protocol Tunneling (L2PT) is not supported on an 802.1Q tunnel port that is mapped for BGP EVPN Layer 2 fabric network extension.
- BGP EVPN Layer 3 IP Routing function is not supported on those VLANs that are mapped to 802.1Q tunnel port.
- Overlapping MAC entry between multiple C-VLANs under a single EVPN-mapped S-VLAN is not supported.

Information About Q-in-VNI

Enterprise campus, data centers, and service provider networks are often required to become a carrier network and provide transparent Layer 2 bridging between statically assigned physical Layer 2 trunk interfaces. Such

networks have specific requirements for VLAN IDs and the number of VLANs to be supported. The VLAN ranges required by different customers in the same service provider network might overlap, and traffic of customers through the infrastructure might be mixed. Assigning a unique range of VLAN IDs to each customer would restrict the customer configurations and could easily exceed the VLAN limit (4094) specified by the IEEE 802.1Q standard.

Using the Q-in-VNI feature, service providers can use a single VLAN (S-VLAN) to support customers who have multiple VLANs (C-VLAN). Each customer's VLAN IDs are preserved, and traffic from different customers is segregated within the service provider network, though they appear to be in the same VLAN. Deploying IEEE 802.1Q tunneling expands the VLAN space by using a VLAN-in-VLAN hierarchy and retagging the tagged packets. A port configured to support IEEE 802.1Q tunneling is called a tunnel port. When tunneling is configured, a tunnel port is assigned to a VLAN ID that is dedicated to tunneling. Each customer is provided with a unique service provider VLAN ID that supports all the VLANs of the customer.

Q-in-VNI in a BGP EVPN VXLAN Fabric

Using the Q-in-VNI feature, a service provider can provide Layer 2 overlay services by mapping the S-VLAN to the Layer 2 VNI. This allows the service providers to address their business customers' Layer 2 connectivity requirements with BGP EVPN VXLAN between the campus sites or a data center.

Enterprise customers can also deploy the Q-in-VNI feature within a single site by mapping the traffic from multiple Layer 2 segments into a specific S-VLAN with EVPN EVI enabled, and with the following criteria:

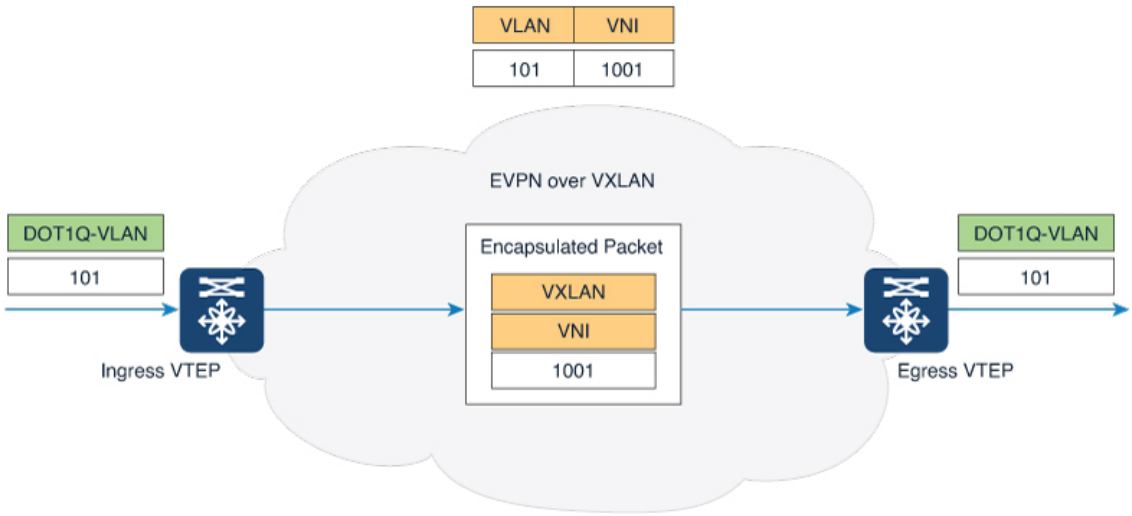
- The site is bounded by the number of L2VNI overlay segments that are supported by a specific Cisco Industrial Ethernet switches.
- VLAN segments are symmetric across the fabric edges.



Note When the Q-in-VNI Layer 2 overlay service with the S-VLAN mapped to an EVPN Instance (also known as MAC VRF) is deployed, the end host MAC routes (RT2) belonging to all the C-VLANs are maintained in a single bridge table corresponding to the S-VLAN.

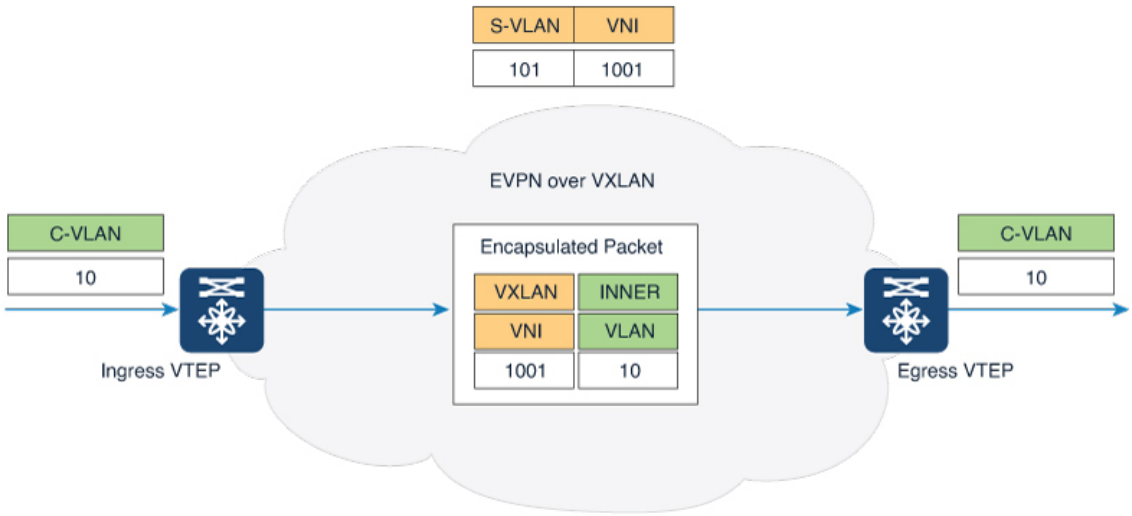
In a BGP EVPN VXLAN fabric with Layer 2 interfaces that have trunk port configuration (Figure), the ingress VTEP strips the IEEE 802.1Q tag and encapsulates a Layer 2 packet with a VXLAN header and forwards the packet to the destination. At the egress VTEP, the packet is decapsulated and L2VNI is mapped to the corresponding VLAN. If the egress port is a trunk port, the corresponding VLAN ID is populated in the IEEE 802.1Q header, and the packet is sent out of the fabric.

Figure 11: EVPN VXLAN Fabric with Trunk Mode on the Access Port



When Q-in-VNI is configured (example topology shown below), customer traffic from C-VLAN with a VLAN ID of 10 is forwarded to the EVPN VXLAN overlay network. The ingress VTEP port in the overlay network is configured as a Q-in-VNI port with a provider VLAN 101 and a unique Layer 2 VNI of 1001. When a packet enters the Q-in-VNI tunnel port on the edge device, it is encapsulated with an outer VXLAN header containing the VNI 1001 (the original inner header with a VLAN 10 is retained). At the Egress VTEP, the packet is forwarded to the correct Q-in-VNI port, based on the matching provider VLAN 101 that is derived from Layer 2 VNI. At the outbound tunnel port, the packet is transmitted with the original C-VLAN tag.

Figure 12: Q-in-VNI over Layer 2 BGP EVPN VXLAN on 802.1Q Tunnel Mode port



How to Configure Q-in-VNI in a BGP EVPN VXLAN Fabric

Configure the access interface for Q-in-VNI tunneling.

Before you begin

Ensure that the Layer 2 overlay is configured for the S-VLAN.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *interface-name* **Example:** `Device(config)# interface GigabitEthernet1/0/1`	Enters interface configuration mode for the interface to be configured as a tunnel port. This should be the edge port on the VTEP that connects to the interface of the Layer 2 device with a trunk port configuration.
Step 4	**switchport access vlan** *vlan-id* **Example:** `Device(config-if)# switchport access vlan 101`	Specifies the S-VLAN that is mapped to the L2VNI.
Step 5	**switchport mode dot1q-tunnel** **Example:** `Device(config-if)# switchport mode dot1q-tunnel`	Sets the interface as an IEEE 802.1Q tunnel port.
Step 6	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Example: Configuring Q-in-VNI in a BGP EVPN VXLAN Fabric

Note Before you enable Q-in-VNI on the interface, ensure that the EVPN VXLAN Layer 2 overlay network is configured. For an example configuration of an EVPN VXLAN Layer 2 overlay, refer to the "Configuring EVPN VXLAN Layer 2 Overlay Network" chapter.

The following example shows how to configure an interface as a tunnel port to the enable Q-in-VNI feature in an EVPN VXLAN Layer 2 overlay network. In this configuration, the VLAN ID for the customer connected to Gigabit Ethernet interface 24 on stack member 1 is VLAN 101.

```
l2vpn evpn instance 101 vlan-based
 encapsulation vxlan
 replication-type static
!
vlan configuration 101        <--- S-VLAN mapped to VNI
 member evpn-instance 101 vni 1001
!
interface nve1
 no ip address
 source-interface Loopback1
 host-reachability protocol bgp
 member vni 1001 mcast-group 225.0.0.101
!
interface GigabitEthernet1/0/1

 switchport access vlan 101   <--- S-VLAN
 switchport mode dot1q-tunnel
 no cdp enable
!
```

Additional References for EVPN VXLAN Layer 2 Overlay with Q-in-VNI

Related Documents

Related Topic	Document Title
IEEE 802.1Q tunneling feature on Cisco Catalyst IE9300 Rugged Series Switches	"Configuring IEEE 802.1Q Tunneling" chapter in the *Layer 2 Configuration Guide, Cisco IOS XE* for the product.



CHAPTER 5

Configuring EVPN VXLAN Layer 3 Overlay Network

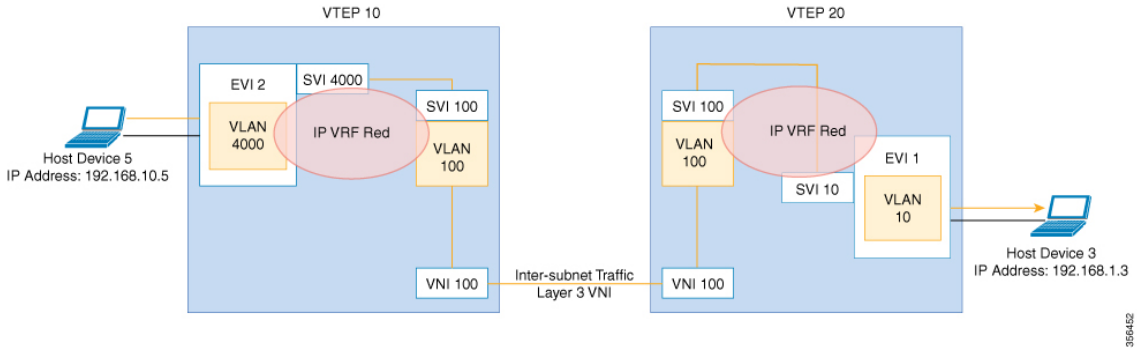
- Information About EVPN VXLAN Layer 3 Overlay Network, on page 83
- How to Configure EVPN VXLAN Layer 3 Overlay Network, on page 84
- Configuration Examples for EVPN VXLAN Layer 3 Overlay Network, on page 94
- Verifying EVPN VXLAN Layer 3 Overlay Network, on page 101

Information About EVPN VXLAN Layer 3 Overlay Network

An EVPN VXLAN Layer 3 overlay network allows host devices in different Layer 2 networks to send Layer 3 or routed traffic to each other. The network forwards the routed traffic using a Layer 3 virtual network instance (VNI) and an IP VRF.

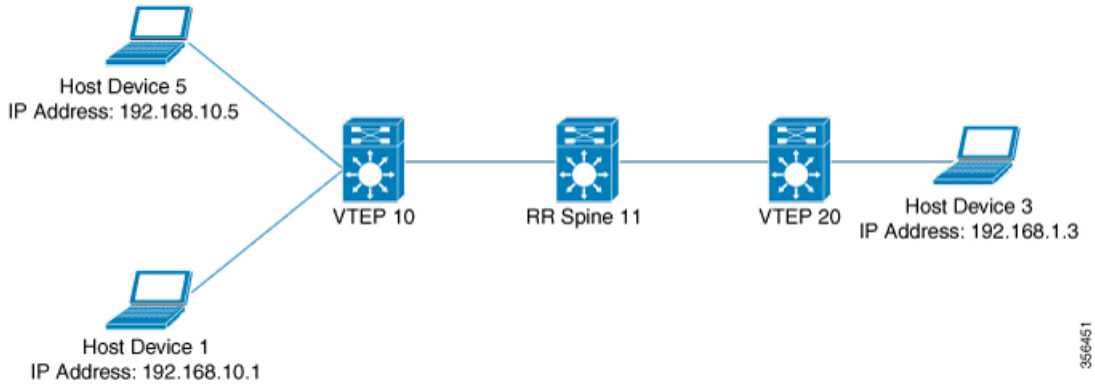
This module provides information only about how to configure a Layer 3 overlay network. You can also configure both Layer 2 and Layer 3 overlay networks together to enable integrated routing and bridging (IRB). For more information about IRB, see *Configuring EVPN VXLAN Integrated Routing and Bridging* module.

The following figure shows the movement of traffic in an EVPN VXLAN Layer 3 overlay network using a Layer 3 VNI:



How to Configure EVPN VXLAN Layer 3 Overlay Network

The following figure shows a sample topology of an EVPN VXLAN Network. Host device 3 and host device 5 are part of different subnets. The network forwards traffic from host device 1 to host device 3 using a Layer 3 VNI and an IP VRF.



Note In a two-VTEP topology, a spine switch is not mandatory. For information about configuration of spine switches in an EVPN VXLAN network, see *Configuring Spine Switches in a BGP EVPN VXLAN Fabric* module.

Perform the following set of procedures to configure an EVPN VXLAN Layer 3 overlay network:

Configuring an IP VRF on a VTEP

To configure an IP VRF on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition Green`	Enters the VRF configuration mode for the specified VRF instance.

	Command or Action	Purpose
Step 4	**rd** *vpn-route-distinguisher* **Example:** `Device(config-vrf)# rd 100:1`	Specifies the route distinguisher for the VRF instance.
Step 5	**address-family ipv4** [**multicast** | **unicast**] **Example:** `Device(config-vrf)# address-family ipv4`	Enters the IPv4 address family configuration mode.

Step 6	**route-target** {**export** | **import** | **both**} *route-target-ext-community*
	Example:
	`Device(config-vrf-af)# route-target export 100:1`
	Example:
	`Device(config-vrf-af)# route-target import 100:1`
		Creates a list of import, export, or both import and export route target communities for the specified VRF.
		Enter either an autonomous system number and an arbitrary number (xxx:y), or an IP address and an arbitrary number (A.B.C.D:y).

Step 7	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **stitching**
	Example:
	`Device(config-vrf-af)# route-target export 100:1 stitching`
	Example:
	`Device(config-vrf-af)# route-target import 100:1 stitching`
		Configures importing, exporting, or both importing and exporting of EVPN route target communities for the VRF.

Step 8	**exit-address-family** **Example:** `Device(config-vrf-af)# exit-address-family`	Exits VRF address family configuration mode and enters VRF configuration mode.
Step 9	**address-family ipv6** [**multicast** | **unicast**] **Example:** `Device(config-vrf)# address-family ipv6`	Enters the IPv6 address family configuration mode.

Step 10	**route-target** {**export** | **import** | **both**} *route-target-ext-community*
	Example:
	`Device(config-vrf-af)# route-target export 100:1`
	Example:
	`Device(config-vrf-af)# route-target import 100:1`
		Creates a list of import, export, or both import and export route target communities for the specified VRF.
		Enter either an autonomous system number and an arbitrary number (xxx:y), or an IP address and an arbitrary number (A.B.C.D:y).

	Command or Action	Purpose

Step 11	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **stitching**
	Example:
	`Device(config-vrf-af)# route-target export 100:1 stitching`
	Example:
	`Device(config-vrf-af)# route-target import 100:1 stitching`
		Configures importing, exporting, or both importing and exporting of VXLAN route target communities for the VRF.

Step 12	**exit-address-family** **Example:** `Device(config-vrf-af)# exit-address-family`	Exits VRF address family configuration mode and enters VRF configuration mode.
Step 13	**end** **Example:** `Device(config-vrf)# end`	Returns to privileged EXEC mode.

Configuring the Core-facing VLAN on a VTEP

To configure the core-facing VLAN on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan configuration** *vlan-id* **Example:** `Device(config)# vlan configuration 11`	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 4	**member vni** *l3-vni-number* **Example:** `Device(config-vlan)# member vni 5000`	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 3 VNI.
Step 5	**end** **Example:**	Returns to privileged EXEC mode.

	Command or Action	Purpose
	`Device(config-vlan)# end`	

Configuring Access-facing VLAN on a VTEP

To configure the access-facing VLAN on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *interface-name* **Example:** `Device(config)# interface GigabitEthernet1/0/1`	Enters interface configuration mode for the specified interface.
Step 4	**switchport access vlan** *vlan-id* **Example:** `Device(config-if)# switchport access vlan 40`	Configures the interface as a static-access port of the specified VLAN. Interface can also be configured as a trunk interface, if required.
Step 5	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring Switch Virtual Interface for the Core-facing VLAN

To configure an SVI for the core-facing VLAN on the VTEP:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *vlan-id* **Example:** `Device(config)# interface vlan 11`	Enters interface configuration mode for the specified VLAN.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding Green`	Configures the SVI for the VLAN.
Step 5	**ip unnumbered***Loopback-interface* **Example:** `Device(config-if)# ip unnumbered Loopback0`	Enables IP processing on the Loopback interface without assigning an explicit IP address to the interface.
Step 6	**no autostate** **Example:** `Device(config-if)# no autostate`	Disables autostate on the interface. In EVPN deployments, once a VLAN is used for a core-facing SVI, it should not be allowed in any trunk. For a core-facing SVI to function properly, the **no autostate** command must be configured under the SVI.
Step 7	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring the Switch Virtual Interface for the Access-facing VLANs

To configure the SVI for the access-facing VLAN on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.

	Command or Action	Purpose
Step 3	**interface vlan** *vlan-id* **Example:** `Device(config)# interface vlan 40`	Enters interface configuration mode for the specified VLAN.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding Green`	Configures the SVI for the VLAN.
Step 5	**ip address***ip-address* **Example:** `Device(config-if)# ip address 192.168.10.100 255.255.255.0`	Configures the IP address of the SVI.
Step 6	**mac-address***mac-address-value* **Example:** `Device(config-if)# mac-address aabb.cc01.f100`	(Optional) Manually sets the MAC address for the VLAN interface.
Step 7	**exit** **Example:** `Device(config-if)# exit`	Returns to global configuration mode.
Step 8	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring the Loopback Interface on a VTEP

To configure the loopback interface on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *loopback-interface-id* **Example:**	Enters interface configuration mode for the specified Loopback interface.

	Command or Action	Purpose
	`Device(config)# interface Loopback0`	
Step 4	**ip address** *ipv4-address* **Example:** `Device(config-if)# ip address 10.12.11.11`	Configures the IP address for the Loopback interface.
Step 5	**ip pim sparse mode** **Example:** `Device(config-if)# ip pim sparse mode`	Enables Protocol Independent Multicast (PIM) sparse mode on the Loopback interface.
Step 6	**end** **Example:** `Device(config-vlan)# end`	Returns to privileged EXEC mode.

Configuring the NVE Interface on a VTEP

To add a VNI member to the NVE interface of a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *nve-interface-id* **Example:** `Device(config)# interface nve1`	Defines the interface to be configured as a trunk, and enters interface configuration mode.
Step 4	**no ip address** **Example:** `Device(config-if)# no ip address`	Disables IP processing on the interface by removing its IP address.
Step 5	**source-interface** *loopback-interface-id* **Example:** `Device(config-if)# source-interface loopback0`	Sets the IP address of the specified loopback interface as the source IP address.
Step 6	**host-reachability protocol bgp** **Example:**	Configures BGP as the host-reachability protocol on the interface.

	Command or Action	Purpose
	`Device(config-if)# host-reachability protocol bgp`	**Note** You must configure the host-reachability protocol on the interface. If you do not execute this step, the VXLAN tunnel defaults to static VXLAN tunnel, which is currently not supported on the Cisco Industrial Ethernet switches.
Step 7	**member vni** *layer2-vni-id* {**ingress-replication** [**local-routing**] | **mcast-group** *multicast-group-address* **Example:** `Device(config-if)# member vni 10000 mcast-group 227.0.0.1`	Associates the Layer 2 VNI member with the NVE. The specified replication type must match the replication type that is configured globally or for the specific EVPN instance. Use **mcast-group** keyword for static replication and **ingress-replication** keyword for ingress replication. Use the **local-routing** keyword only when you need to configure route type 2 (RT 2) to route type 5 (RT 5) reorigination on the centralized gateway (CGW) VTEP.
Step 8	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring BGP with EVPN and VRF Address Families on a VTEP

To configure BGP on a VTEP with EVPN and VRF address families and a spine switch as the neighbor, perform these steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 1`	Enables a BGP routing process, assigns it an autonomous system number, and enters router configuration mode.

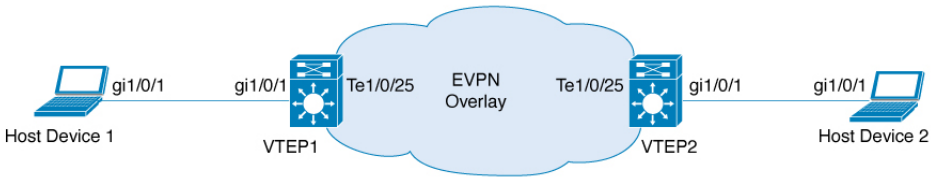
	Command or Action	Purpose
Step 4	**bgp log-neighbor-changes** **Example:** `Device(config-router)# bgp log-neighbor-changes`	(Optional) Enables the generation of logging messages when the status of a BGP neighbor changes. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 5	**bgp update-delay** *time-period* **Example:** `Device(config-router)# bgp update-delay 1`	(Optional) Sets the maximum initial delay period before sending the first update. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 6	**bgp graceful-restart** **Example:** `Device(config-router)# bgp graceful-restart`	(Optional) Enables the BGP graceful restart capability for all BGP neighbors. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 7	**no bgp default ipv4-unicast** **Example:** `Device(config-router)# no bgp default ipv4-unicast`	(Optional) Disables default IPv4 unicast address family for BGP peering session establishment. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 8	**neighbor** *ip-address* **remote-as** *number* **Example:** `Device(config-router)# neighbor 10.11.11.11 remote-as 1`	Defines multiprotocol-BGP neighbors. Under each neighbor, define the Layer 2 Virtual Private Network (L2VPN) EVPN configuration. Use the IP address of the spine switch as the neighbor IP address.
Step 9	**neighbor** {*ip-address* | *group-name*} **update-source** *interface* **Example:** `Device(config-router)# neighbor 10.11.11.11 update-source Loopback0`	Configures update source. Update source can be configured per neighbor or per peer-group. Use the IP address of the spine switch as the neighbor IP address.
Step 10	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Specifies the L2VPN address family and enters address family configuration mode.
Step 11	**neighbor** *ip-address* **activate** **Example:** `Device(config-router-af)# neighbor 10.11.11.11 activate`	Enables the exchange information from a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 12	**neighbor** *ip-address* **send-community** [**both** | **extended** | **standard**]	Specifies the communities attribute sent to a BGP neighbor.

	Command or Action	Purpose
	Example: `Device(config-router-af)# neighbor 10.11.11.11 send-community both`	Use the IP address of the spine switch as the neighbor IP address.
Step 13	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 14	**address-family ipv4 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv4 vrf green`	Specifies the IPv4 address family and enters address family configuration mode.
Step 15	**advertise l2vpn evpn***limit valueallow-suppressed* **Example:** `Device(config-router-af)# advertise l2vpn evpn <1-2147483647> allow-suppressed`	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric. <1-2147483647> Defines upper limit on advertise/export prefixes without hogging memory. <allow-suppressed> also advertises suppressed routes.
Step 16	**redistribute connected** **Example:** `Device(config-router-af)# redistribute connected`	Redistributes connected routes to BGP.
Step 17	**redistribute static** **Example:** `Device(config-router-af)# redistribute static`	Redistributes static routes to BGP.
Step 18	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 19	**address-family ipv6 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv6 vrf green`	Specifies the IPv6 address family and enters address family configuration mode.
Step 20	**advertise l2vpn evpn***limit valueallow-suppressed* **Example:**	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric. <1-2147483647> Defines upper limit on advertise/export prefixes without hogging memory.

	Command or Action	Purpose
	`Device(config-router-af)# advertise l2vpn evpn <1-2147483647> allow-suppressed`	<allow-suppressed> also advertises suppressed routes.
Step 21	**redistribute connected** **Example:** `Device(config-router-af)# redistribute connected`	Redistributes connected routes to BGP.
Step 22	**redistribute static** **Example:** `Device(config-router-af)# redistribute static`	Redistributes static routes to BGP.
Step 23	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 24	**end** **Example:** `Device(config-router)# end`	Returns to privileged EXEC mode.

Configuration Examples for EVPN VXLAN Layer 3 Overlay Network

This section provides an example for configuring an EVPN VXLAN Layer 3 overlay network. This example shows a sample configuration for a VXLAN network with two VTEPs, VTEP 1 and VTEP 2, connected to perform routing.



Note In a two-VTEP topology, a spine switch is not mandatory. For information about configuration of spine switches in an EVPN VXLAN network, see *Configuring Spine Switches in a BGP EVPN VXLAN Fabric* module.

Table 8: Configuration Example for a VXLAN Network with Two VTEPs Connected to Perform Routing

VTEP 1	VTEP 2

VTEP 1	VTEP 2
Device# **show running-config** ! hostname Device ! ! vrf definition green rd 103:2 ! address-family ipv4 route-target export 103:2 route-target import 104:2 route-target export 103:2 stitching route-target import 104:2 stitching exit-address-family ! address-family ipv6 route-target export 103:2 route-target import 104:2 route-target export 103:2 stitching route-target import 104:2 stitching exit-address-family ! ip multicast-routing ipv6 unicast-routing ! ! system mtu 9150 ! vlan configuration 200 member vni 5000 ! ! interface Loopback0 ip address 10.1.1.10 255.255.255.255 ip pim sparse-mode ! interface Loopback13 description demo only (for rt5 distribution) vrf forwarding green ip address 10.1.13.13 255.255.255.0 ! interface GigabitEthernet1/0/1 description access interface switchport access vlan 201 switchport mode access ! ! interface Te1/0/25 description core-underlay-interface no switchport ip address 172.16.1.29 255.255.255.0 ip pim sparse-mode ! ! interface Vlan200 description core svi for l3vni vrf forwarding green ip unnumbered Loopback0 ipv6 enable no autostate ! interface Vlan201	Device# **show running-config** ! hostname Device ! ! vrf definition green rd 104:2 ! address-family ipv4 route-target export 104:2 route-target import 103:2 route-target export 104:2 stitching route-target import 103:2 stitching exit-address-family ! address-family ipv6 route-target export 104:2 route-target import 103:2 route-target export 104:2 stitching route-target import 103:2 stitching exit-address-family ! ip multicast-routing ipv6 unicast-routing ! ! system mtu 9150 ! vlan configuration 200 member vni 5000 ! ! interface Loopback0 ip address 10.2.2.20 255.255.255.255 ip pim sparse-mode ! interface Loopback14 description demo only (for rt5 distribution) vrf forwarding green ip address 10.1.14.14 255.255.255.0 ! interface GigabitEthernet1/0/1 description access interface switchport access vlan 202 switchport mode access ! ! interface Te1/0/25 description core-underlay-interface no switchport ip address 172.16.1.30 255.255.255.0 ip pim sparse-mode ! ! interface Vlan200 description core svi for l3vni vrf forwarding green ip unnumbered Loopback0 ipv6 enable no autostate ! interface Vlan202

VTEP 1	VTEP 2
description access-svi vrf forwarding green ip address 192.168.1.201 255.255.255.0 ipv6 address 2001:DB8:201::201/64 ipv6 enable ! interface nve10 no ip address source-interface Loopback0 host-reachability protocol bgp member vni 5000 vrf green ! router ospf 1 router-id 10.1.1.10 network 10.1.1.0 0.0.0.255 area 0 network 172.16.1.0 0.0.0.255 area 0 ! router bgp 10 bgp router-id interface Loopback0 bgp log-neighbor-changes bgp update-delay 1 no bgp default ipv4-unicast neighbor 10.2.2.20 remote-as 10 neighbor 10.2.2.20 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 10.2.2.20 activate neighbor 10.2.2.20 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 10.1.1.10 ! ! end	description access-svi vrf forwarding green ip address 192.168.2.202 255.255.255.0 ipv6 address 2001:DB8:202::202/64 ipv6 enable ! interface nve10 no ip address source-interface Loopback0 host-reachability protocol bgp member vni 5000 vrf green ! router ospf 1 router-id 10.2.2.20 network 10.2.2.0 0.0.0.255 area 0 network 172.16.1.0 0.0.0.255 area 0 ! router bgp 10 bgp router-id interface Loopback0 bgp log-neighbor-changes bgp update-delay 1 no bgp default ipv4-unicast neighbor 10.1.1.10 remote-as 10 neighbor 10.1.1.10 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 10.1.1.10 activate neighbor 10.1.1.10 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 10.1.1.10 ! ! end

The following examples provide outputs for **show** commands on VTEP 1 and VTEP 2 in the topology configured above.

- #unique_68 unique_68_Connect_42_section_zll_qxs_nkb
- #unique_68 unique_68_Connect_42_section_zwz_pxs_nkb
- #unique_68 unique_68_Connect_42_section_y3n_pxs_nkb
- #unique_68 unique_68_Connect_42_section_jyv_4xs_nkb

show nve peers

VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve10      5000     L3CP 10.2.2.20        380e.4d9b.6a4a 5000       UP  A/M/4 00:38:37
nve10      5000     L3CP 10.2.2.20        380e.4d9b.6a4a 5000       UP  A/-/6 00:03:16
```

VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Device# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve10      5000     L3CP 10.1.1.10        a0f8.4910.bce2 5000       UP  A/-/4 00:38:53
nve10      5000     L3CP 10.1.1.10        a0f8.4910.bce2 5000       UP  A/M/6 00:38:53
```

show bgp l2vpn evpn all

VTEP 1

The following example shows the output for the **show bgp l2vpn evpn all** all command on VTEP 1:

```
Device# show bgp l2vpn evpn all
BGP table version is 26, local router ID is 10.1.1.10
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 103:2 (default for vrf green)
 *>   [5][103:2][0][24][10.1.13.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][103:2][0][24][192.168.1.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][103:2][0][64][2001:DB8:201::]/29
                      ::                       0         32768 ?
Route Distinguisher: 104:2
 *>i  [5][104:2][0][24][10.1.14.0]/17
                      10.2.2.20                0    100      0 ?
 *>i  [5][104:2][0][24][192.168.2.0]/17
                      10.2.2.20                0    100      0 ?
 *>i  [5][104:2][0][64][2001:DB8:202::]/29
                      10.2.2.20                0    100      0 ?
```

VTEP 2

The following example shows the output for the **show bgp l2vpn evpn all** command on VTEP 2:

```
Device# show bgp l2vpn evpn all
BGP table version is 12, local router ID is 10.2.2.20
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 103:2
 *>i  [5][103:2][0][24][10.1.13.0]/17
                      10.1.1.10                0    100      0 ?
 *>i  [5][103:2][0][24][192.168.1.0]/17
                      10.1.1.10                0    100      0 ?
 *>i  [5][103:2][0][64][2001:DB8:201::]/29
                      10.1.1.10                0    100      0 ?
Route Distinguisher: 104:2 (default for vrf green)
 *>   [5][104:2][0][24][10.1.14.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][104:2][0][24][192.168.2.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][104:2][0][64][2001:DB8:202::]/29
     Network          Next Hop            Metric LocPrf Weight Path
                      ::                       0         32768 ?
```

show ip route vrf

VTEP 1

The following example shows the output for the **show ip route vrf** command on VTEP 1:

```
Device# show ip route vrf green
Routing Table: green
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, 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, 3 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Loopback13
L        10.1.13.13/32 is directly connected, Loopback13
B        10.1.14.0/24 [200/0] via 10.2.2.20, 00:42:01, Vlan200
      192.168.1.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Vlan201
L        192.168.1.201/32 is directly connected, Vlan201
B     192.168.2.0/24 [200/0] via 10.2.2.20, 00:06:00, Vlan200
```

VTEP 2

The following example shows the output for the **show ip route vrf** command on VTEP 2:

```
Device# show ip route vrf green
Routing Table: green
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, 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, 3 subnets, 2 masks
B        10.1.13.0/24 [200/0] via 10.1.1.10, 00:42:38, Vlan200
C        10.1.14.0/24 is directly connected, Loopback14
L        10.1.14.14/32 is directly connected, Loopback14
B     192.168.1.0/24 [200/0] via 10.1.1.10, 00:42:38, Vlan200
      192.168.2.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.2.0/24 is directly connected, Vlan202
L        192.168.2.202/32 is directly connected, Vlan202
```

show platform software fed switch active matm mactable vlan

VTEP 1

The following example shows the output for the **show platform software fed switch active matm mactable vlan 200** command on VTEP 1:



Note The MAC address of the peer's core SVI interface must be present in the core VLAN.

```
Device# show platform software fed switch active matm macTable vlan 200
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle            diHandle             *a_time  *e_time  ports

-------------------------------------------------------------------------------------------------------------------------
200    a0f8.4910.bce2      0x8002      0  19880     64  0x7f5d8503fd48      0x7f5d852b6d28
      0x0                 0x5234                       0        0  Vlan200

200    380e.4d9b.6a4a   0x1000001      0      0     64  0x7f5d85117598      0x7f5d85110f78
      0x7f5d851b9648      0x0                          0        0  RLOC 10.2.2.20 adj_id 22


Total Mac number of addresses:: 2
```

VTEP 2

The following example shows the output for the **show platform software fed switch active matm mactable vlan 200** command on VTEP 2:

Note The MAC address of the peer's core SVI interface must be present in the core VLAN.

```
Device# show platform software fed switch active matm macTable vlan 200
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle           diHandle             *a_time  *e_time  ports

-------------------------------------------------------------------------------------------------------------------------------
200    380e.4d9b.6a4a      0x8002      0  42949     64  0x7f40e15fd308      0x7f40e15f49d8
      0x0                 0x0                        0        0  Vlan200

200    a0f8.4910.bce2   0x1000001      0      0     64  0x7f40e193c478      0x7f40e1938168
      0x7f40e1937bf8      0x0                        0        0  RLOC 10.1.1.10 adj_id 86


Total Mac number of addresses:: 2
```

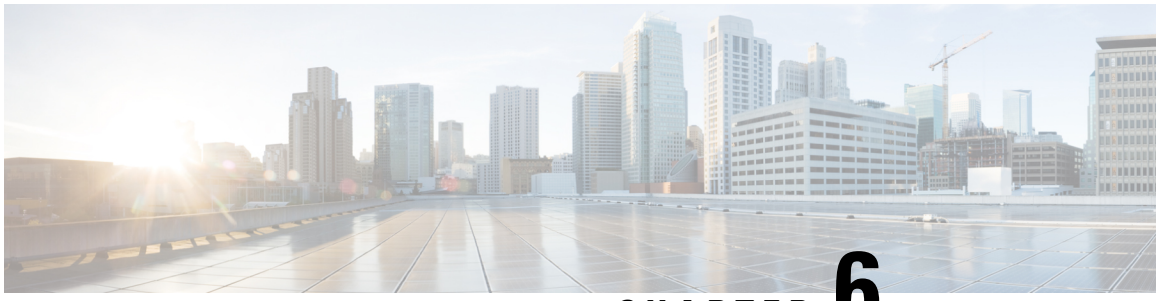
Verifying EVPN VXLAN Layer 3 Overlay Network

The following table lists the **show** commands that are used to verify a Layer 3 VXLAN overlay network:

Table 9: Commands to Verify EVPN VXLAN Layer 3 Overlay Network

Command	Purpose
show nve vni	Displays information about VXLAN network identifier members associated with an NVE interface.
show nve vni *vni-id* **detail**	Displays detailed NVE interface state information for a VXLAN network identifier member.
show nve peers	Displays NVE interface state information for peer leaf switches.
show mac address-table vlan *vlan-id*	Displays MAC addresses for a VLAN.
show platform software fed switch active matm macTable vlan *vlan-id*	Displays MAC addresses for a VLAN from MAC address table manager database for Forwarding Engine Driver (FED).
show ip route vrf *vrf-name*	Displays the IP routing table associated with a specific VRF.
show ip cef vrf *vrf-name*	Displays entries in the Cisco Express Forwarding (CEF) table associated with a VRF.
show arp vrf *vrf-name*	Displays entries in the Address Resolution Protocol (ARP) table associated with a VRF.
show bgp l2vpn evpn route-type 5	Displays BGP information for route type 5 of Layer 2 VPN EVPN address family.

Command	Purpose
show bgp l2vpn evpn all	Displays all BGP information for L2VPN EVPN address family.



CHAPTER 6

Configuring EVPN VXLAN Integrated Routing and Bridging

- Restrictions for EVPN VXLAN Integrated Routing and Bridging, on page 103
- Integrated Routing and Bridging, on page 104
- How to Configure EVPN VXLAN Integrated Routing and Bridging, on page 108
- Verifying EVPN VXLAN Integrated Routing and Bridging, on page 123
- Configuration Examples for EVPN VXLAN Integrated Routing and Bridging, on page 123

Restrictions for EVPN VXLAN Integrated Routing and Bridging

This section provides restrictions for both EVPN VXLAN distributed anycast gateway and centralized default gateway functionalities that are used to enable integrated routing and bridging (IRB).

Layer 2 or Layer 3 access control lists (ACLs) are not supported for the received or transmitted packets that have VXLAN encapsulation.

EVPN VXLAN Distributed Anycast Gateway

The same subnet mask and IP address must be configured on all the switch virtual interfaces (SVIs) that act as a distributed anycast gateway (DAG).

EVPN VXLAN Centralized Default Gateway

- Only one centralized gateway (CGW) leaf switch or VTEP is supported in an EVPN VXLAN network. To achieve physical redundancy, configure Cisco Stackwise Virtual on the leaf switches. For more information, see *Configuring Cisco StackWise Virtual* module of the *High Availability Configuration Guide*.
- HSRP and VRRP are not supported for the EVPN VXLAN Layer 2 overlay networks when you use centralized gateway.
- Reorigination of route type 2 host routes to route type 5 routes is supported.

Integrated Routing and Bridging

EVPN VXLAN supports Integrated Routing and Bridging (IRB) functionality which allows the VTEPs in a VXLAN network to forward both Layer 2 (bridged) and Layer 3 (routed) traffic. When a VTEP forwards Layer 2 traffic, it is said to be performing bridging. Similarly, when a VTEP forwards Layer 3 traffic, it is said to be performing routing. The traffic between different subnets is forwarded through the VXLAN gateways. IRB is implemented in two ways:

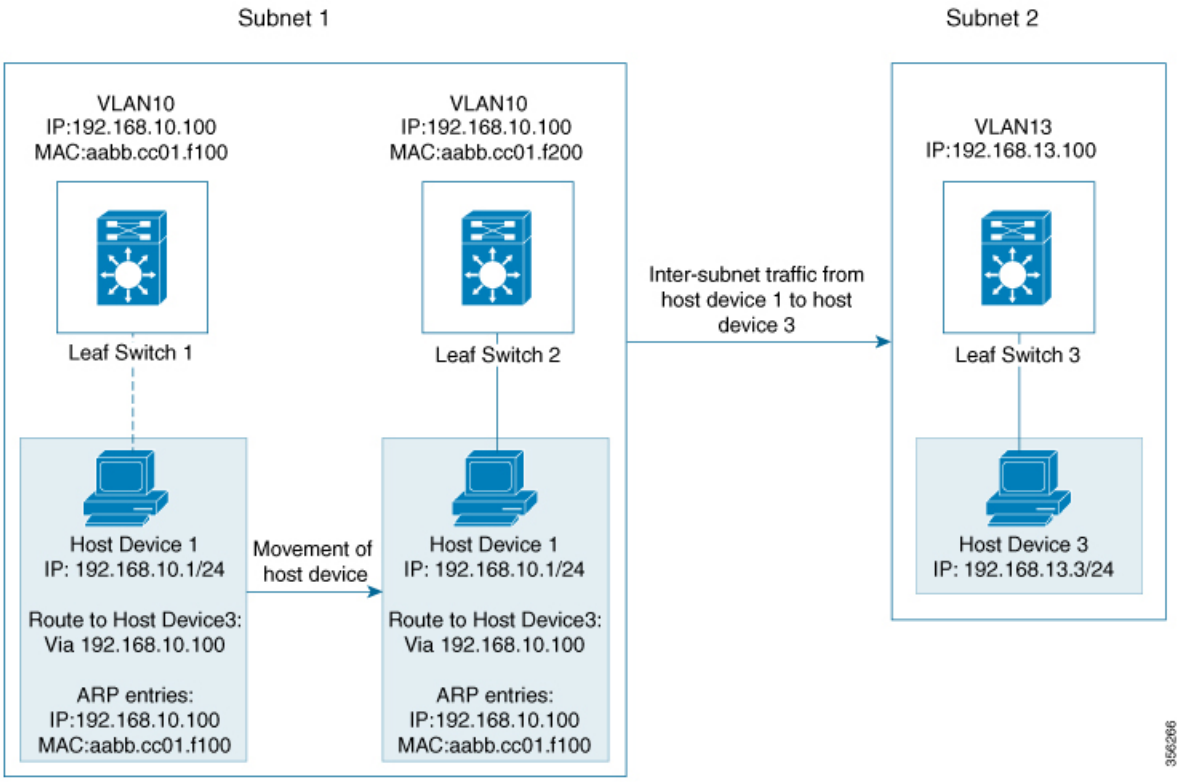
- Asymmetric IRB
- Symmetric IRB

EVPN VXLAN Distributed Anycast Gateway

Distributed anycast gateway is a default gateway addressing mechanism in a BGP EVPN VXLAN fabric. The feature enables the use of the same gateway IP and MAC address across all the VTEPs in an EVPN VXLAN network. This ensures that every VTEP functions as the default gateway for the workloads directly connected to it. The feature facilitates flexible workload placement, host mobility, and optimal traffic forwarding across the BGP EVPN VXLAN fabric.

The scenario shown in the following figure depicts a distributed gateway. Subnet 1 contains two leaf switches, leaf switch 1 and leaf switch 2, acting together as a distributed default gateway for VLAN 10. Host device 1 is connected to leaf switch 1 and needs to send traffic to host device 3, which is in a different subnet. When host device 1 tries to send traffic outside of subnet 1, the traffic goes through the configured gateway on leaf switch 1. Host device 1 registers the Address Resolution Protocol (ARP) entries of the gateway VLAN MAC and IP addresses on leaf switch 1.

Figure 13: Distributed Gateway Topology



When multiple VTEPs act together as one single distributed default gateway for the same VLAN, the VLAN IP address remains the same across all of them. This IP address becomes the gateway IP address for any host device in the VLAN that tries to reach an IP address outside its subnet. But, each VTEP retains its own MAC address.

In the preceding figure, consider the scenario where host device 1 moves from leaf switch 1 to leaf switch 2. The host device remains within the same network and still maintains the same ARP entries for gateway MAC and IP addresses. But the MAC addresses of the VLAN interfaces on leaf switch 2 and leaf switch 1 are different. This results in a MAC address mismatch between the ARP entry and the VLAN on leaf switch 2. As a result, any traffic that host device 1 tries to send outside of Subnet 1 is either lost or continuously flooded as unknown unicast. EVPN VXLAN distributed anycast gateway feature prevents this traffic loss by ensuring that all the VTEPs have the same gateway MAC and IP addresses.

Manual MAC address configuration and MAC aliasing are the two methods used to maintain the same MAC address across all VTEPs and configure distributed anycast gateway.

Manual MAC Address Configuration

Manual MAC address configuration is the conventional method of enabling distributed anycast gateway in an EVPN VXLAN network. In this method, you manually configure the same MAC address on the Layer 2 VNI VLAN SVI on all the VTEPs in the network. You must configure the same MAC address on all the VTEPs in the same Layer 2 VNI.

Note The VLAN SVIs on all the leaf switches must already share the same gateway IP address.

MAC Aliasing

MAC aliasing for distributed anycast gateway removes the need to configure the same MAC address explicitly on the VLAN interfaces of every VTEP. MAC aliasing allows the VTEPs to advertise their VLAN MAC addresses as the gateway MAC addresses to all the other VTEPs in the network. The VTEPs in the network store the advertised MAC address as a gateway MAC address provided their VLAN IP address matches with the gateway IP address.

MAC aliasing in an EVPN VXLAN network is configured by enabling the default gateway advertisement (using **default-gateway advertise** command) on all the VTEPs.

Note If the same MAC address is assigned to all the VTEPs and the **default-gateway advertise** command is configured, only the unique gateway MAC addresses are imported into the default gateway table.

EVPN VXLAN Centralized Default Gateway

In this scenario, the network has a CGW VTEP that performs the Layer 3 gateway function for all the Layer 2 VNIs. All the other VTEPs in the network perform only bridging. The CGW VTEP acts as the Layer 3 gateway and performs routing for the inter-subnet VXLAN traffic.

The CGW VTEP advertises the SVI MAC-IP route for a particular VXLAN-enabled VLAN to all other Layer 2 VTEPs that have the same Layer 2 VNI configured. This allows the VTEPs to import and install the remote SVI MAC-IP route as a VXLAN Layer 3 gateway address. A host device uses the address of an SVI in the same VLAN on the CGW VTEP as its gateway address. Configure the SVI for the Layer 2 VNI VLAN only on the CGW VTEP. Do not configure the SVI (for the respective Layer 2 VNI VLAN) on any other VTEP in the network that acts as a Layer 2 VTEP.

When a host device connected to a Layer 2 VTEP sends traffic to a different subnet, the traffic is bridged from the Layer 2 VTEP to the CGW VTEP. The CGW VTEP then routes the traffic to the destination subnet. The destination subnet can be another VXLAN-enabled VLAN or an external route.

If the CGW VTEP needs to route the traffic between 2 VXLAN-enabled VLANs, then configure the CGW on the same VTEP for both VLANs. In other words, configure the SVI on the same VTEP for both VLANs.

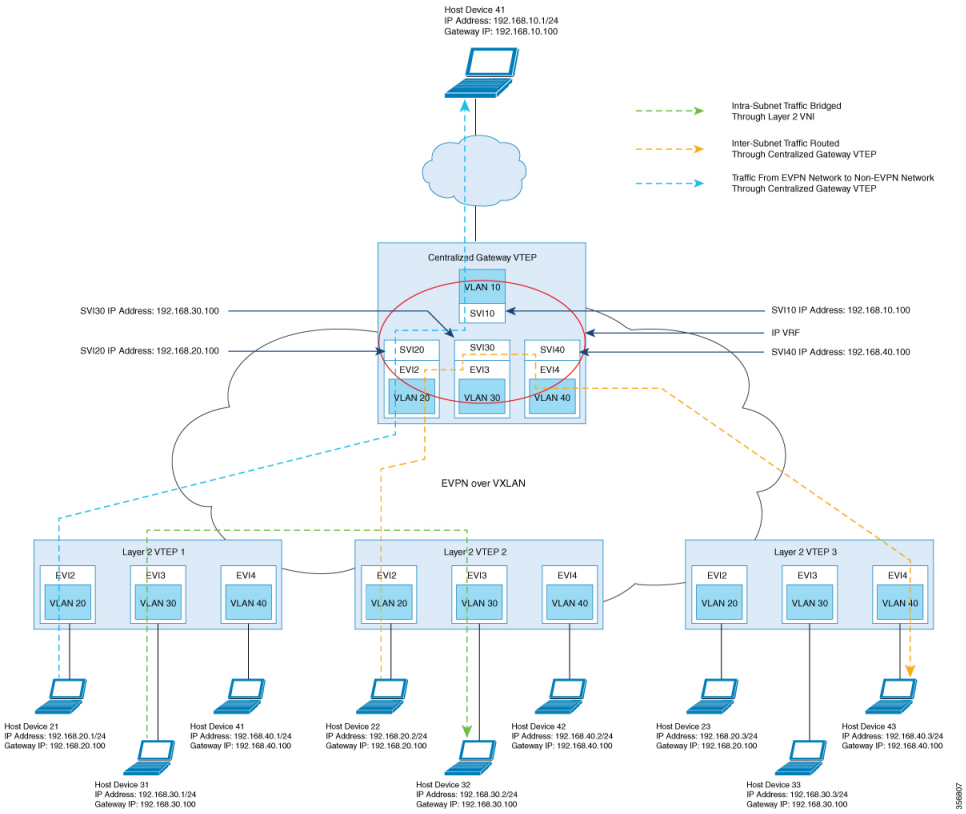
We recommend that you configure a centralized default gateway in an EVPN VXLAN network if:

- You require a boundary between Layer 2 and Layer 3 segments at the border of the BGP EVPN VXLAN fabric.
- The inter-subnet traffic is subjected to a firewall inspection or any policy on a centralized plane.

Note Toggling between DAG and CGW on an SVI in a VLAN disrupts the traffic for that VLAN.

The following image shows an EVPN VXLAN network with centralized default gateway configured:

Figure 14: An EVPN VXLAN Network with Centralized Default Gateway



Default Gateway MAC Address Assignment

When leaf switches import gateway addresses, it can result in a conflict if the SVI of a leaf switch has the same IP and MAC address as the imported addresses. To avoid this conflict, the SVI MAC-IP routes are tagged with the Default Gateway Extended Community attribute. The attribute helps the receiving leaf switches to distinguish the MAC-IP routes of the SVIs from the MAC-IP routes of the host devices. When a leaf switch receives a route tagged with the attribute, it results in one of the following scenarios:

- If the leaf switch does not have a local SVI for the same MAC VRF, it installs the route only as a remote MAC route. The leaf switch implements centralized gateway functionality in this scenario.
- If the leaf switch has a local SVI with a matching IP address but different MAC address, it installs the MAC route as a route that points to the local SVI. The leaf switch implements MAC aliasing for distributed anycast gateway in this scenario.
- If the leaf switch has an SVI with no matching IP address, it invalidates the MAC-IP route and issues an error.

Route Type 2 to Route Type 5 Host Route Reorigination for Overlay Networks with a Centralized Gateway (Asymmetric IRB)

Route type 2 (RT 2) to Route type 5 (RT 5) reorigination allows a Layer 2 and Layer 3 VXLAN network to import IP addresses from another Layer 2-only VXLAN network. In this deployment model, a VXLAN gateway VTEP functions as both the CGW and the Layer 3 VXLAN gateway. You can use the gateway VTEP to connect a Layer 2-only VXLAN network with a Layer 2 and Layer 3 VXLAN network. The gateway VTEP uses the CGW and asymmetric IRB forwarding functionalities in the EVPN segment of the Layer 2-only network. The VTEP can reoriginate any RT 2 host routes, that it learns from the Layer 2 EVPN segment (or MAC VRF), as RT 5 routes in the Layer 3 EVPN segment (or IP VRF). The VTEP then forwards the reoriginated RT 5 routes to host devices in the Layer 2 and Layer 3 VXLAN network.



Note The Layer 3 EVPN segment can be one of the following:

- An IP VRF-only segment (see *Configuring EVPN VXLAN Layer 3 Overlay Network* module for more information)
- Distributed anycast gateway (DAG) overlay fabric

To enable RT 2 to RT 5 host route reorigination on the CGW VTEP, ensure that you run the **reoriginate route-type5** and **member vni** *layer2-vni-id* **ingress-replication local-routing** commands on that VTEP.



Note The CGW VTEP does not install the reoriginated RT 5 route on its local IP VRF.

How to Configure EVPN VXLAN Integrated Routing and Bridging

To configure EVPN VXLAN Integrated Routing and Bridging, you need to configure EVPN VXLAN Layer 2 and Layer 3 overlay networks, and enable the gateways in the VXLAN network.

Configuring EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway

To enable IRB in an EVPN VXLAN network using distributed anycast gateway, perform the following set of procedures:

- Configure Layer 2 VPN EVPN on the VTEPs.

  Enable distributed anycast gateway for the VXLAN network when you configure Layer 2 VPN.
- Configure the core-facing and access-facing VLANs on the VTEPs.
- Configure switch virtual interface (SVI) for the core-facing VLAN on the VTEPs.
- Configure SVI for the access-facing VLAN on the VTEPs.
- Configure the IP VRF on the VTEPs.

- Configure the Loopback interface on the VTEPs.
- Configure the Network Virtualization Endpoint (NVE) interface on the VTEPs.
- Configure BGP with EVPN address family on the VTEPs.

Configuring Layer 2 VPN EVPN on a VTEP

To configure the Layer 2 VPN EVPN parameters on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**l2vpn evpn** **Example:** `Device(config)# l2vpn evpn`	Enters EVPN configuration mode.
Step 4	**replication-type** {**ingress** | **static**} **Example:** `Device(config-evpn)# replication-type static`	Configures the Layer 2 VPN EVPN replication type. **Note** Configure the Layer 2 VPN EVPN replication type as static, if multicast is enabled in the underlay network for EVPN BUM traffic. When the Layer 2 VPN EVPN replication type is configured as static, the IMET route is not advertised and forwarding of BUM traffic relies on underlay multicast being configured on each VTEP.
Step 5	**router-id** *loopback-interface-id* **Example:** `Device(config-evpn)# router-id loopback 0`	Specifies the interface that will supply the IP addresses to be used in auto-generating route distinguishers.
Step 6	**default-gateway advertise** **Example:** `Device(config-evpn)# default-gateway advertise`	(Optional) Enables default gateway advertisement on the switch. To configure distributed anycast gateway in a VXLAN network using MAC aliasing, enable default gateway advertisement on all the leaf switches in the network.

	Command or Action	Purpose
		This command is applicable in integrated routing and bridging (IRB) scenarios where Layer 2 and Layer 3 VNIs coexist in a VRF. Refer to *Configuring EVPN VXLAN Integrated Routing and Bridging* module for more details. This command is mandatory only if the same MAC address is not manually configured on all the access SVIs. **Note** Use the **default-gateway advertise** {**enable** | **disable**} command in EVPN instnace configuration mode to override the global default gateway advertisement settings and enable or disable it for a specific EVPN instance.
Step 7	**logging peer state** **Example:** `Device(config-evpn)# logging peer state`	(Optional) Displays syslog message when the first route is received or the last route is withdrawn from a given remote VTEP.
Step 8	**mac duplication limit** *limit-number* **time** *time-limit* **Example:** `Device(config-evpn)# mac duplication limit 20 time 5`	(Optional) Changes parameters for detecting duplicate MAC addresses.
Step 9	**ip duplication limit** *limit-number* **time** *time-limit* **Example:** `Device(config-evpn)# ip duplication limit 20 time 5`	(Optional) Changes parameters for detecting duplicate IP addresses.
Step 10	**route-target auto vni** **Example:** `Device(config-evpn)# route-target auto vni`	(Optional) Specifies to use VNI instead of EVPN instance number to auto-generate route target.
Step 11	**exit** **Example:** `Device(config-evpn)# exit`	Exits EVPN configuration mode and enters global configuration mode.
Step 12	**l2vpn evpn instance** *evpn-instance-number* **vlan-based** **Example:** `Device(config)# l2vpn evpn instance 1 vlan-based`	Configures a VLAN based EVPN instance in Layer 2 VPN configuration mode. An EVPN instance needs to be explicitly configured only when something needs to be

	Command or Action	Purpose
		configured per EVPN instance such as a route target.
Step 13	**encapsulation vxlan** **Example:** `Device(config-evpn-evi)# encapsulation vxlan`	(Optional) Defines the encapsulation format as VXLAN. The encapsulation format is VXLAN by default.
Step 14	**replication-type** {**ingress** | **static**} **Example:** `Device(config-evpn-evi)# replication-type ingress`	(Optional) Sets the replication type for the EVPN instance. In case a global replication type has already been configured, this overrides the global setting.
Step 15	**default-gateway advertise** {**enable** | **disable**} **Example:** `Device(config-evpn-evi)# default-gateway advertise disable`	(Optional) Enables or disables the default gateway advertisement for the EVPN instance. In case default gateway advertisement has already been globally configured, this overrides the global setting. This command is mandatory only if the same MAC address is not manually configured on all the access SVIs. To configure distributed anycast gateway in a VXLAN network using MAC aliasing, enable default gateway advertisement on all the leaf switches in the network.
Step 16	**ip local-learning** {**enable** | **disable**} **Example:** `Device(config-evpn-evi)# ip local-learning disable`	(Optional) Enables or disables local IP address learning for the specified EVPN instance. In case IP address learning has already been globally configured, this overrides the global setting.
Step 17	**re-originate route-type5** **Example:** `Device(config-evpn-evi)# re-originate route-type5`	(Optional) Enables the centralized gateway (CGW) VTEP to re-originate the route-type 2 (RT 2) host routes from a Layer 2 VTEP as route-type 5 (RT 5) network routes into a Layer 3 overlay network.
Step 18	**no auto-route-target** **Example:** `Device(config-evpn-evi)# no auto-route-target`	(Optional) Disables auto generation of route targets.
Step 19	**rd** *rd-value* **Example:** `Device(config-evpn-evi)# rd 65000:100`	(Optional) Configures a route distinguisher manually.

	Command or Action	Purpose
Step 20	**route-target** {**import** | **export** | **both**} *rt-value* **Example:** `Device(config-evpn-evi)# route-target both 65000:100`	(Optional) Configures route targets manually. **Note** Configure route targets manually if the auto-generated route target values (ASN:EVI or ASN:VNI) are different between the VTEPs.
Step 21	**end** **Example:** `Device(config-evpn-evi)# end`	Returns to privileged EXEC mode.

Configuring an IP VRF on a VTEP

To configure an IP VRF on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition Green`	Enters the VRF configuration mode for the specified VRF instance.
Step 4	**rd** *vpn-route-distinguisher* **Example:** `Device(config-vrf)# rd 100:1`	Specifies the route distinguisher for the VRF instance.
Step 5	**address-family ipv4** [**multicast** | **unicast**] **Example:** `Device(config-vrf)# address-family ipv4`	Enters the IPv4 address family configuration mode.
Step 6	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **Example:** `Device(config-vrf-af)# route-target export 100:1`	Creates a list of import, export, or both import and export route target communities for the specified VRF. Enter either an autonomous system number and an arbitrary number (xxx:y), or an IP address and an arbitrary number (A.B.C.D:y).

	Command or Action	Purpose
	Example: `Device(config-vrf-af)# route-target import 100:1`	

Step 7	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **stitching**
	Example:
	`Device(config-vrf-af)# route-target export 100:1 stitching`
	Example:
	`Device(config-vrf-af)# route-target import 100:1 stitching`
		Configures importing, exporting, or both importing and exporting of EVPN route target communities for the VRF.

Step 8	**exit-address-family** **Example:** `Device(config-vrf-af)# exit-address-family`	Exits VRF address family configuration mode and enters VRF configuration mode.
Step 9	**address-family ipv6** [**multicast** | **unicast**] **Example:** `Device(config-vrf)# address-family ipv6`	Enters the IPv6 address family configuration mode.

Step 10	**route-target** {**export** | **import** | **both**} *route-target-ext-community*
	Example:
	`Device(config-vrf-af)# route-target export 100:1`
	Example:
	`Device(config-vrf-af)# route-target import 100:1`
		Creates a list of import, export, or both import and export route target communities for the specified VRF.
		Enter either an autonomous system number and an arbitrary number (xxx:y), or an IP address and an arbitrary number (A.B.C.D:y).

Step 11	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **stitching**
	Example:
	`Device(config-vrf-af)# route-target export 100:1 stitching`
	Example:
	`Device(config-vrf-af)# route-target import 100:1 stitching`
		Configures importing, exporting, or both importing and exporting of VXLAN route target communities for the VRF.

Step 12	**exit-address-family** **Example:** `Device(config-vrf-af)# exit-address-family`	Exits VRF address family configuration mode and enters VRF configuration mode.

	Command or Action	Purpose
Step 13	**end** **Example:** `Device(config-vrf)# end`	Returns to privileged EXEC mode.

Configuring the Loopback Interface on a VTEP

To configure the loopback interface on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *loopback-interface-id* **Example:** `Device(config)# interface Loopback0`	Enters interface configuration mode for the specified Loopback interface.
Step 4	**ip address** *ipv4-address* **Example:** `Device(config-if)# ip address 10.12.11.11`	Configures the IP address for the Loopback interface.
Step 5	**ip pim sparse mode** **Example:** `Device(config-if)# ip pim sparse mode`	Enables Protocol Independent Multicast (PIM) sparse mode on the Loopback interface.
Step 6	**end** **Example:** `Device(config-vlan)# end`	Returns to privileged EXEC mode.

Configuring Core-facing and Access-facing VLANs on a VTEP

To configure the core-facing and access-facing VLANs on a VTEP and enable IRB in the EVPN VXLAN network, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan configuration** *vlan-id* **Example:** `Device(config)# vlan configuration 201`	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 4	**member evpn-instance** *evpn-instance-id* **vni** *l2-vni-number* **Example:** `Device(config-vlan)# member evpn-instance 1 vni 6000`	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 2 VNI.
Step 5	**exit** **Example:** `Device(config-vlan)# exit`	Returns to global configuration mode.
Step 6	**vlan configuration** *vlan-id* **Example:** `Device(config)# vlan configuration 202`	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 7	**member evpn-instance** *evpn-instance-id* **vni** *l2-vni-number* **Example:** `Device(config-vlan)# member evpn-instance 2 vni 7000`	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 2 VNI.
Step 8	**exit** **Example:** `Device(config-vlan)# exit`	Returns to global configuration mode.
Step 9	**vlan configuration** *vlan-id* **Example:** `Device(config)# vlan configuration 200`	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 10	**member vni** *l3-vni-number* **Example:**	Adds EVPN instance as a member of the VLAN configuration.

	Command or Action	Purpose
	`Device(config-vlan)# member vni 5000`	The VNI here is used as a Layer 3 VNI.
Step 11	**exit** **Example:** `Device(config-vlan)# exit`	Returns to global configuration mode.
Step 12	**end** **Example:** `Device(config-vlan)# end`	Returns to privileged EXEC mode.

Configuring Switch Virtual Interface for the Core-facing VLAN on a VTEP

To configure an SVI for the core-facing VLAN on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *vlan-id* **Example:** `Device(config)# interface vlan 200`	Enters interface configuration mode for the specified VLAN.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding Green`	Configures the SVI for the VLAN.
Step 5	**ip unnumbered***Loopback-interface* **Example:** `Device(config-if)# ip unnumbered Loopback0`	Enables IP processing on the Loopback interface without assigning an explicit IP address to the interface.
Step 6	**no autostate** **Example:** `Device(config-if)# no autostate`	Disables autostate on the interface. In EVPN deployments, once a VLAN is used for a core-facing SVI, it should not be allowed in any trunk. For a core-facing SVI to function properly, the **no autostate** command must be configured under the SVI.

	Command or Action	Purpose
Step 7	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring Switch Virtual Interface for the Access-facing VLANs on a VTEP

To configure SVIs for the access-facing VLANs on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *vlan-id* **Example:** `Device(config)# interface vlan 202`	Enters interface configuration mode for the specified VLAN.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding Green`	Configures the SVI for the VLAN.
Step 5	**ip address** *gateway-ip-address* **Example:** `Device(config-if)# ip address 192.168.10.1 255.255.255.0`	Configures the gateway IP address for the access SVI. Configure the same gateway IP address for this SVI on all the other VTEPs.
Step 6	**mac-address** *mac-address-value* **Example:** `Device(config-if)# mac-address aabb.cc01.f100`	(Optional) Manually sets the MAC address for the VLAN interface. To configure distributed anycast gateway in a VXLAN network using manual MAC configuration, configure the same MAC address on the corresponding Layer 2 VNI SVIs on all the VTEPs in a VXLAN network.
Step 7	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring the NVE Interface on a VTEP

To add Layer 2 and Layer 3 VNI members to the NVE interface of a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *nve-interface-id* **Example:** `Device(config)# interface nve1`	Defines the interface to be configured as a trunk, and enters interface configuration mode.
Step 4	**no ip address** **Example:** `Device(config-if)# no ip address`	Disables IP processing on the interface by removing its IP address.
Step 5	**source-interface** *loopback-interface-id* **Example:** `Device(config-if)# source-interface loopback0`	Sets the IP address of the specified loopback interface as the source IP address.
Step 6	**host-reachability protocol bgp** **Example:** `Device(config-if)# host-reachability protocol bgp`	Configures BGP as the host-reachability protocol on the interface. **Note** You must configure the host reachability protocol on the interface. If you do not execute this step, the VXLAN tunnel defaults to static VXLAN tunnel, which is currently not supported on the Cisco Industrial Ethernet switches.

Step 7	**member vni** *layer2-vni-id* {**ingress-replication** [**local-routing**] | **mcast-group** *multicast-group-address* }
	Example:
	`Device(config-if)# member vni 6000 mcast-group 227.0.0.1`
	`Device(config-if)# member vni 7000 mcast-group 227.0.0.1`
		Associates the Layer 2 VNI member with the NVE.
		The specified replication type must match the replication type that is configured globally or for the specific EVPN instance. Use **mcast-group** keyword for static replication and **ingress-replication** keyword for ingress replication.

	Command or Action	Purpose
		Use the **local-routing** keyword only when you need to configure route type 2 (RT 2) to route type 5 (RT 5) reorigination on the centralized gateway (CGW) VTEP.
Step 8	**member vni** *layer3-vni-id* **vrf** *vrf-name* **Example:** `Device(config-if)# member vni 5000 vrf Green`	Associates the Layer 3 VNI member with the NVE.
Step 9	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring BGP with EVPN and VRF Address Families on a VTEP

To configure BGP on a VTEP with EVPN and VRF address families and a spine switch as the neighbor, perform these steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 1`	Enables a BGP routing process, assigns it an autonomous system number, and enters router configuration mode.
Step 4	**bgp log-neighbor-changes** **Example:** `Device(config-router)# bgp log-neighbor-changes`	(Optional) Enables the generation of logging messages when the status of a BGP neighbor changes. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 5	**bgp update-delay** *time-period* **Example:** `Device(config-router)# bgp update-delay 1`	(Optional) Sets the maximum initial delay period before sending the first update. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.

	Command or Action	Purpose
Step 6	**bgp graceful-restart** **Example:** `Device(config-router)# bgp graceful-restart`	(Optional) Enables the BGP graceful restart capability for all BGP neighbors. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 7	**no bgp default ipv4-unicast** **Example:** `Device(config-router)# no bgp default ipv4-unicast`	(Optional) Disables default IPv4 unicast address family for BGP peering session establishment. For more information, see *Configuring BGP* module of the *IP Routing Configuration Guide*.
Step 8	**neighbor** *ip-address* **remote-as** *number* **Example:** `Device(config-router)# neighbor 10.11.11.11 remote-as 1`	Defines multiprotocol-BGP neighbors. Under each neighbor, define the Layer 2 Virtual Private Network (L2VPN) EVPN configuration. Use the IP address of the spine switch as the neighbor IP address.
Step 9	**neighbor** {*ip-address* | *group-name*} **update-source** *interface* **Example:** `Device(config-router)# neighbor 10.11.11.11 update-source Loopback0`	Configures update source. Update source can be configured per neighbor or per peer-group. Use the IP address of the spine switch as the neighbor IP address.
Step 10	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Specifies the L2VPN address family and enters address family configuration mode.
Step 11	**neighbor** *ip-address* **activate** **Example:** `Device(config-router-af)# neighbor 10.11.11.11 activate`	Enables the exchange information from a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 12	**neighbor** *ip-address* **send-community** [**both** | **extended** | **standard**] **Example:** `Device(config-router-af)# neighbor 10.11.11.11 send-community both`	Specifies the communities attribute sent to a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 13	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 14	**address-family ipv4 vrf** *vrf-name* **Example:**	Specifies the IPv4 address family and enters address family configuration mode.

	Command or Action	Purpose
	`Device(config-router)# address-family ipv4 vrf green`	
Step 15	**advertise l2vpn evpn***limit valueallow-suppressed* **Example:** `Device(config-router-af)# advertise l2vpn evpn <1-2147483647> allow-suppressed`	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric. <1-2147483647> Defines upper limit on advertise/export prefixes without hogging memory. <allow-suppressed> also advertises suppressed routes.
Step 16	**redistribute connected** **Example:** `Device(config-router-af)# redistribute connected`	Redistributes connected routes to BGP.
Step 17	**redistribute static** **Example:** `Device(config-router-af)# redistribute static`	Redistributes static routes to BGP.
Step 18	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 19	**address-family ipv6 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv6 vrf green`	Specifies the IPv6 address family and enters address family configuration mode.
Step 20	**advertise l2vpn evpn***limit valueallow-suppressed* **Example:** `Device(config-router-af)# advertise l2vpn evpn <1-2147483647> allow-suppressed`	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric. <1-2147483647> Defines upper limit on advertise/export prefixes without hogging memory. <allow-suppressed> also advertises suppressed routes.
Step 21	**redistribute connected** **Example:** `Device(config-router-af)# redistribute connected`	Redistributes connected routes to BGP.
Step 22	**redistribute static** **Example:**	Redistributes static routes to BGP.

	Command or Action	Purpose
	`Device(config-router-af)# redistribute static`	
Step 23	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 24	**end** **Example:** `Device(config-router)# end`	Returns to privileged EXEC mode.

Configuring EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway

This section provides information about how to configure EVPN VXLAN IRB using centralized default gateway. Centralized default gateway implements asymmetric IRB in the EVPN VXLAN network. To enable IRB using centralized default gateway in an EVPN VXLAN network, perform these steps:

- On the centralized gateway VTEP, perform these procedures:
  - Configuring Layer 2 VPN EVPN on a VTEP, on page 17
  - Configuring the Loopback Interface on a VTEP, on page 21
  - Configuring Core-facing and Access-facing VLANs on a VTEP, on page 114
  - Configuring Switch Virtual Interface for the Core-facing VLAN on a VTEP, on page 116
  - Configuring Switch Virtual Interface for the Access-facing VLANs on a VTEP, on page 117
  - Configuring the NVE Interface on a VTEP, on page 118
  - Configuring BGP with EVPN and VRF Address Families on a VTEP, on page 91

Note

- Do not configure a Layer 3 VNI on the centralized gateway VTEP for the specific VRF in which the Layer 2 VNIs are configured.
- In order to set the SVI state to up for the respective EVPN instance, run the **no autostate** command in the VLAN configuration mode for the SVI.
- The centralized gateway leaf switch can be configured either as a single switch or a stack of switches..

- On all the other VTEPs, configure only EVPN VXLAN Layer 2 overlay network. To configure EVPN VXLAN Layer 2 overlay network on the leaf switches, perform all the procedures listed in How to Configure EVPN VXLAN Layer 2 Overlay Network, on page 16.

Note As Layer 2 VTEPs perform only bridging, do not configure the SVIs on the non-centralized-gateway VTEPs.

Verifying EVPN VXLAN Integrated Routing and Bridging

The following sections provide information about how to verify EVPN VXLAN integrated routing and bridging:

Verifying EVPN VXLAN Layer 2 Overlay Network

See Verifying EVPN VXLAN Layer 2 Overlay Network, on page 25 for the list of commands.

Verifying EVPN VXLAN Layer 3 Overlay Network

See Verifying EVPN VXLAN Layer 3 Overlay Network, on page 101 for the list of commands.

Verifying Distributed Anycast Gateway

The following table lists the **show** commands that are used to verify distributed anycast gateway:

Table 10: Commands to Verify Distributed Anycast Gateway

Command	Purpose
show l2vpn evpn default-gateway	Displays the default gateway database.
show l2vpn l2route default-gateway	Displays the list of sent or received default gateway routes.
show mac address-table	Displays the list of MAC addresses received in default gateway routes that are installed as static MAC addresses for an SVI interface.

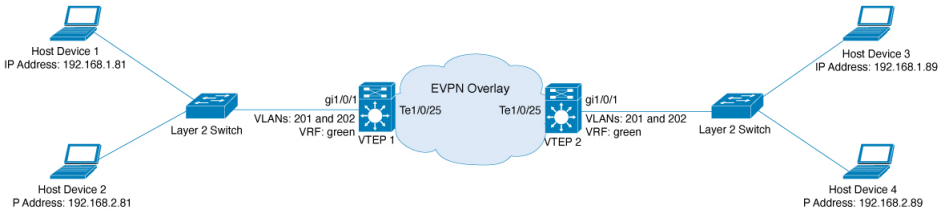
Configuration Examples for EVPN VXLAN Integrated Routing and Bridging

This section provides examples to show how EVPN VXLAN IRB is enabled using distributed anycast gateway and centralized default gateway.

Example: Enabling EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway

This section provides an example to show how to enable EVPN VXLAN IRB using distributed anycast gateway. The following example shows a sample configuration for an EVPN VXLAN network with 2 VTEPs. VTEP 1 and VTEP 2 are connected to perform integrated routing and bridging.

Figure 15: EVPN VXLAN Topology for IRB using Distributed Anycast Gateway



Note In a two-VTEP topology, a spine switch is not mandatory. For information about configuration of spine switches in an EVPN VXLAN network, see *Configuring Spine Switches in a BGP EVPN VXLAN Fabric* module.

Table 11: Configuration Example for an EVPN VXLAN Network with Two VTEPs Connected to Perform Integrated Routing and Bridging Using Distributed Anycast Gateway

VTEP 1	VTEP 2

VTEP 1	VTEP 2

VTEP1# **show running-config**
!
hostname VTEP1
!
vrf definition green
 rd 103:2
 !
 address-family ipv4
  route-target export 103:2
  route-target import 104:2
  route-target export 103:2 stitching
  route-target import 104:2 stitching
 exit-address-family
 !
 address-family ipv6
  route-target export 103:2
  route-target import 104:2
  route-target export 103:2 stitching
  route-target import 104:2 stitching
 exit-address-family
!
ip routing
ip multicast-routing
ipv6 unicast-routing
!
!
l2vpn evpn
 replication-type static
 router-id Loopback0
 default-gateway advertise
!
l2vpn evpn instance 1 vlan-based
 encapsulation vxlan
!
l2vpn evpn instance 2 vlan-based
 encapsulation vxlan
!
!
system mtu 9150
!
vlan configuration 200
 member vni 5000
vlan configuration 201
 member evpn-instance 1 vni 6000
vlan configuration 202
 member evpn-instance 2 vni 7000
!
!
interface Loopback0
 ip address 10.1.1.10 255.255.255.255
 ip pim sparse-mode
!
interface Loopback13
description demo only (for rt5 distribution)
 vrf forwarding green
 ip address 10.1.13.13 255.255.255.0
!
interface GigabitEthernet1/0/1
description access-facing-interface
 switchport trunk allowed vlan 201,202
 switchport mode trunk
!
	VTEP2# **show running-config**
	!
	hostname VTEP2
	!
	vrf definition green
	 rd 104:2
	 !
	 address-family ipv4
	  route-target export 104:2
	  route-target import 103:2
	  route-target export 104:2 stitching
	  route-target import 103:2 stitching
	 exit-address-family
	 !
	 address-family ipv6
	  route-target export 104:2
	  route-target import 103:2
	  route-target export 104:2 stitching
	  route-target import 103:2 stitching
	 exit-address-family
	!
	ip routing
	ip multicast-routing
	ipv6 unicast-routing
	!
	!
	l2vpn evpn
	 replication-type static
	 router-id Loopback0
	 default-gateway advertise
	!
	l2vpn evpn instance 1 vlan-based
	 encapsulation vxlan
	!
	l2vpn evpn instance 2 vlan-based
	 encapsulation vxlan
	!
	!
	system mtu 9150
	!
	vlan configuration 200
	 member vni 5000
	vlan configuration 201
	 member evpn-instance 1 vni 6000
	vlan configuration 202
	 member evpn-instance 2 vni 7000
	!
	!
	interface Loopback0
	 ip address 10.2.2.20 255.255.255.255
	 ip pim sparse-mode
	!
	interface Loopback14
	 description demo only (for rt5 distribution)
	 vrf forwarding green
	 ip address 10.1.14.14 255.255.255.0
	!
	interface GigabitEthernet1/0/1
	description access-facing-interface
	 switchport trunk allowed vlan 201,202
	 switchport mode trunk
	!

VTEP 1	VTEP 2
!	!
interface Te1/0/25	interface Te1/0/25
description core-underlay-interface	description core-underlay-interface
no switchport	no switchport
ip address 172.16.1.29 255.255.255.0	ip address 172.16.1.30 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
!	!
!	!
interface Vlan200	interface Vlan200
description core svi for l3vni	description core svi for l3vni
vrf forwarding green	vrf forwarding green
ip unnumbered Loopback0	ip unnumbered Loopback0
ipv6 enable	ipv6 enable
no autostate	no autostate
!	!
interface Vlan201	interface Vlan201
description vni 6000 default-gateway	description vni 6000 default-gateway
vrf forwarding green	vrf forwarding green
ip address 192.168.1.201 255.255.255.0	ip address 192.168.1.201 255.255.255.0
ipv6 address 2001:DB8:201::201/64	ipv6 address 2001:DB8:201::201/64
ipv6 enable	ipv6 enable
!	!
interface Vlan202	interface Vlan202
description vni 7000 default-gateway	description vni 7000 default-gateway
vrf forwarding green	vrf forwarding green
ip address 192.168.2.202 255.255.255.0	ip address 192.168.2.202 255.255.255.0
ipv6 address 2001:DB8:202::202/64	ipv6 address 2001:DB8:202::202/64
ipv6 enable	ipv6 enable
!	!
!	!
interface nve10	interface nve10
no ip address	no ip address
source-interface Loopback0	source-interface Loopback0
host-reachability protocol bgp	host-reachability protocol bgp
member vni 6000 mcast-group 232.1.1.1	member vni 6000 mcast-group 232.1.1.1
member vni 5000 vrf green	member vni 7000 mcast-group 232.1.1.1
member vni 7000 mcast-group 232.1.1.1	member vni 5000 vrf green
!	!
router ospf 1	router ospf 1
router-id 10.1.1.10	router-id 10.2.2.20
network 10.1.1.0 0.0.0.255 area 0	network 10.2.2.0 0.0.0.255 area 0
network 172.16.1.0 0.0.0.255 area 0	network 172.16.1.0 0.0.0.255 area 0
!	!
router bgp 10	router bgp 10
bgp router-id interface Loopback0	bgp router-id interface Loopback0
bgp log-neighbor-changes	bgp log-neighbor-changes
bgp update-delay 1	bgp update-delay 1
no bgp default ipv4-unicast	no bgp default ipv4-unicast
neighbor 10.2.2.20 remote-as 10	neighbor 10.1.1.10 remote-as 10
neighbor 10.2.2.20 update-source Loopback0	neighbor 10.1.1.10 update-source Loopback0
!	!
address-family ipv4	address-family ipv4
exit-address-family	exit-address-family
!	!
address-family l2vpn evpn	address-family l2vpn evpn
neighbor 10.2.2.20 activate	neighbor 10.1.1.10 activate
neighbor 10.2.2.20 send-community both	neighbor 10.1.1.10 send-community both
exit-address-family	exit-address-family
!	!
address-family ipv4 vrf green	address-family ipv4 vrf green
advertise l2vpn evpn	advertise l2vpn evpn
redistribute connected	redistribute connected

VTEP 1	VTEP 2
redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 10.1.1.10 ! end	redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 10.1.1.10 ! end

The following examples provide outputs for **show** commands on VTEP 1 and VTEP 2 in the topology configured above:

- show nve peers, on page 128
- show l2vpn evpn peers vxlan, on page 129
- show l2vpn evpn evi evpn-instance detail, on page 129
- show l2vpn evpn default-gateway, on page 130
- show bgp l2vpn evpn all, on page 131
- show ip route vrf green, on page 134
- show platform software fed switch active matm mactable vlan, on page 135

show nve peers

VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
VTEP1# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve10      5000     L3CP 10.2.2.20        380e.4d9b.6a4a 5000       UP  A/M/4 01:33:41
nve10      5000     L3CP 10.2.2.20        380e.4d9b.6a4a 5000       UP  A/-/6 00:43:38
nve10      6000     L2CP 10.2.2.20        5              6000       UP   N/A  01:33:41
nve10      7000     L2CP 10.2.2.20        6              7000       UP   N/A  01:33:41
```

VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
VTEP2# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve10      5000     L3CP 10.1.1.10        a0f8.4910.bce2 5000       UP  A/M/4 01:33:55
nve10      5000     L3CP 10.1.1.10        a0f8.4910.bce2 5000       UP  A/-/6 01:14:23
nve10      6000     L2CP 10.1.1.10        7              6000       UP   N/A  01:33:55
nve10      7000     L2CP 10.1.1.10        6              7000       UP   N/A  01:33:55
```

show l2vpn evpn peers vxlan

VTEP 1

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
VTEP1# l2vpn evpn peers vxlan
Interface VNI      Peer-IP                  Num routes eVNI     UP time
--------- -------- -----------------------  ---------- -------- --------
nve10     6000     10.2.2.20                5          6000     01:34:50
nve10     7000     10.2.2.20                6          7000     01:34:50
```

VTEP 2

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
VTEP2# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                  Num routes eVNI     UP time
--------- -------- -----------------------  ---------- -------- --------
nve10     6000     10.1.1.10                7          6000     01:35:23
nve10     7000     10.1.1.10                6          7000     01:35:23
```

show l2vpn evpn evi evpn-instance detail

VTEP 1

The following example shows the output for the **show l2vpn evpn evi** *evpn-instance* **detail** command on VTEP 1:

```
VTEP1# show l2vpn evpn evi 1 detail
EVPN instance:      1 (VLAN Based)
  RD:               10.1.1.10:1 (auto)
  Import-RTs:       10:1
  Export-RTs:       10:1
  Per-EVI Label:    none
  State:            Established
  Replication Type: Static (global)
  Encapsulation:    vxlan
  IP Local Learn:   Enable (global)
  Vlan:             201
    Ethernet-Tag:   0
    State:          Established
    Core If:        Vlan200
    Access If:      Vlan201
    NVE If:         nve10
    RMAC:           a0f8.4910.bce2
    Core Vlan:      200
    L2 VNI:         6000
    L3 VNI:         5000
    VTEP IP:        10.1.1.10
    MCAST IP:       232.1.1.1
    VRF:            green
    IPv4 IRB:       Enabled
    IPv6 IRB:       Enabled
    Pseudoports:
      GigabitEthernet1/0/1 service instance 201
```

VTEP 2

The following example shows the output for the **show l2vpn evpn evi** *evpn-instance* **detail** command on VTEP 2:

```
VTEP2# show l2vpn evpn evi 1 detail
EVPN instance:      1 (VLAN Based)
  RD:               10.2.2.20:1 (auto)
  Import-RTs:       10:1
  Export-RTs:       10:1
  Per-EVI Label:    none
  State:            Established
  Replication Type: Static (global)
  Encapsulation:    vxlan
  IP Local Learn:   Enable (global)
  Vlan:             201
    Ethernet-Tag:   0
    State:          Established
    Core If:        Vlan200
    Access If:      Vlan201
    NVE If:         nve10
    RMAC:           380e.4d9b.6a4a
    Core Vlan:      200
    L2 VNI:         6000
    L3 VNI:         5000
    VTEP IP:        10.2.2.20
    MCAST IP:       232.1.1.1
    VRF:            green
    IPv4 IRB:       Enabled
    IPv6 IRB:       Enabled
    Pseudoports:
      GigabitEthernet1/0/1 service instance 201
```

show l2vpn evpn default-gateway

VTEP 1

The following example shows the output for the **show l2vpn evpn default-gateway** command on VTEP 1:

```
VTEP1# show l2vpn evpn default-gateway
Valid Default Gateway Address   EVI   VLAN  MAC Address    Source
----- ------------------------- ----- ----- -------------- --------------------
  Y   192.168.1.201             1     201   a0f8.4910.bccc Vl201
  Y   192.168.1.201             1     201   380e.4d9b.6a48 10.2.2.20
  Y   2001:DB8:201::201         1     201   a0f8.4910.bccc Vl201
  Y   2001:DB8:201::201         1     201   380e.4d9b.6a48 10.2.2.20
  Y   192.168.2.202             2     202   a0f8.4910.bcc2 Vl202
  Y   192.168.2.202             2     202   380e.4d9b.6a42 10.2.2.20
  Y   2001:DB8:202::202         2     202   a0f8.4910.bcc2 Vl202
  Y   2001:DB8:202::202         2     202   380e.4d9b.6a42 10.2.2.20
```

VTEP 2

The following example shows the output for the **show l2vpn evpn default-gateway** command on VTEP 2:

```
VTEP2# show l2vpn evpn default-gateway
Valid Default Gateway Address   EVI   VLAN  MAC Address    Source
----- ------------------------ ----- ----- -------------- --------------------
  Y   192.168.1.201            1     201   380e.4d9b.6a48 Vl201
  Y   192.168.1.201            1     201   a0f8.4910.bccc 10.1.1.10
  Y   2001:DB8:201::201        1     201   380e.4d9b.6a48 Vl201
  Y   2001:DB8:201::201        1     201   a0f8.4910.bccc 10.1.1.10
  Y   192.168.2.202            2     202   380e.4d9b.6a42 Vl202
  Y   192.168.2.202            2     202   a0f8.4910.bcc2 10.1.1.10
  Y   2001:DB8:202::202        2     202   380e.4d9b.6a42 Vl202
  Y   2001:DB8:202::202        2     202   a0f8.4910.bcc2 10.1.1.10
```

show bgp l2vpn evpn all

VTEP 1

The following example shows the output for the **show bgp l2vpn evpn all** command on VTEP 1:

```
VTEP1# show bgp l2vpn evpn all
BGP table version is 705, local router ID is 10.1.1.10
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.10:1
 *>i  [2][10.1.1.10:1][0][48][0018736C56C3][0][*]/20
                      10.2.2.20                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][0018736C56C3][32][192.168.1.89]/24
                      10.2.2.20                0    100      0 ?
 *>   [2][10.1.1.10:1][0][48][0059DC50AE01][0][*]/20
                      ::                                    32768 ?
 *>   [2][10.1.1.10:1][0][48][0059DC50AE4C][0][*]/20
                      ::                                    32768 ?
 *>   [2][10.1.1.10:1][0][48][0059DC50AE4C][32][192.168.1.81]/24
                      ::                                    32768 ?
 *>   [2][10.1.1.10:1][0][48][0059DC50AE4C][128][2001:DB8:201::81]/36
                      ::                                    32768 ?
 *>   [2][10.1.1.10:1][0][48][0059DC50AE4C][128][FE80::259:DCFF:FE50:AE4C]/36
                      ::                                    32768 ?
 *>i  [2][10.1.1.10:1][0][48][380E4D9B6A48][32][192.168.1.201]/24
                      10.2.2.20                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][380E4D9B6A48][128][2001:DB8:201::201]/36
                      10.2.2.20                0    100      0 ?
 *>   [2][10.1.1.10:1][0][48][A0F84910BCCC][32][192.168.1.201]/24
                      ::                                    32768 ?
 *>   [2][10.1.1.10:1][0][48][A0F84910BCCC][128][2001:DB8:201::201]/36
                      ::                                    32768 ?
Route Distinguisher: 10.1.1.10:2
 *>i  [2][10.1.1.10:2][0][48][0018736C5681][0][*]/20
                      10.2.2.20                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][0018736C56C2][0][*]/20
                      10.2.2.20                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][0018736C56C2][32][192.168.2.89]/24
                      10.2.2.20                0    100      0 ?
 *>   [2][10.1.1.10:2][0][48][0059DC50AE01][0][*]/20
                      ::                                    32768 ?
 *>   [2][10.1.1.10:2][0][48][0059DC50AE42][0][*]/20
```

```
                    ::                                    32768 ?
 *>   [2][10.1.1.10:2][0][48][0059DC50AE42][32][192.168.2.81]/24
                    ::                                    32768 ?
 *>i  [2][10.1.1.10:2][0][48][380E4D9B6A42][32][192.168.2.202]/24
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][380E4D9B6A42][128][2001:DB8:202::202]/36
                    10.2.2.20                0    100      0 ?
 *>   [2][10.1.1.10:2][0][48][A0F84910BCC2][32][192.168.2.202]/24
                    ::                                    32768 ?
 *>   [2][10.1.1.10:2][0][48][A0F84910BCC2][128][2001:DB8:202::202]/36
                    ::                                    32768 ?
Route Distinguisher: 10.2.2.20:1
 *>i  [2][10.2.2.20:1][0][48][0018736C56C3][0][*]/20
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][0018736C56C3][32][192.168.1.89]/24
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][380E4D9B6A48][32][192.168.1.201]/24
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][380E4D9B6A48][128][2001:DB8:201::201]/36
                    10.2.2.20                0    100      0 ?
Route Distinguisher: 10.2.2.20:2
 *>i  [2][10.2.2.20:2][0][48][0018736C5681][0][*]/20
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:2][0][48][0018736C56C2][0][*]/20
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:2][0][48][0018736C56C2][32][192.168.2.89]/24
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:2][0][48][380E4D9B6A42][32][192.168.2.202]/24
                    10.2.2.20                0    100      0 ?
 *>i  [2][10.2.2.20:2][0][48][380E4D9B6A42][128][2001:DB8:202::202]/36
                    10.2.2.20                0    100      0 ?
Route Distinguisher: 103:2 (default for vrf green)
 *>   [5][103:2][0][24][10.1.13.0]/17
                    0.0.0.0                  0         32768 ?
 *>   [5][103:2][0][24][192.168.1.0]/17
                    0.0.0.0                  0         32768 ?
 *>   [5][103:2][0][24][192.168.2.0]/17
                    0.0.0.0                  0         32768 ?
 *>   [5][103:2][0][64][2001:DB8:201::]/29
                    ::                       0         32768 ?
 *>   [5][103:2][0][64][2001:DB8:202::]/29
                    ::                       0         32768 ?
Route Distinguisher: 104:2
 *>i  [5][104:2][0][24][10.1.14.0]/17
                    10.2.2.20                0    100      0 ?
 *>i  [5][104:2][0][24][192.168.1.0]/17
                    10.2.2.20                0    100      0 ?
 *>i  [5][104:2][0][24][192.168.2.0]/17
                    10.2.2.20                0    100      0 ?
 *>i  [5][104:2][0][64][2001:DB8:201::]/29
                    10.2.2.20                0    100      0 ?
 *>i  [5][104:2][0][64][2001:DB8:202::]/29
                    10.2.2.20                0    100      0 ?
```

VTEP 2

The following example shows the output for the **show bgp l2vpn evpn all** command on VTEP 2:

```
VTEP2# show bgp l2vpn evpn all
BGP table version is 584, local router ID is 10.2.2.20
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 10.1.1.10:1
 *>i  [2][10.1.1.10:1][0][48][0059DC50AE01][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][0059DC50AE4C][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][0059DC50AE4C][32][192.168.1.81]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][0059DC50AE4C][128][2001:DB8:201::81]/36
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][0059DC50AE4C][128][FE80::259:DCFF:FE50:AE4C]/36
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][A0F84910BCCC][32][192.168.1.201]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:1][0][48][A0F84910BCCC][128][2001:DB8:201::201]/36
                      10.1.1.10                0    100      0 ?
Route Distinguisher: 10.1.1.10:2
 *>i  [2][10.1.1.10:2][0][48][0059DC50AE01][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][0059DC50AE42][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][0059DC50AE42][32][192.168.2.81]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][A0F84910BCC2][32][192.168.2.202]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.1.1.10:2][0][48][A0F84910BCC2][128][2001:DB8:202::202]/36
                      10.1.1.10                0    100      0 ?
Route Distinguisher: 10.2.2.20:1
 *>   [2][10.2.2.20:1][0][48][0018736C56C3][0][*]/20
                      ::                                    32768 ?
 *>   [2][10.2.2.20:1][0][48][0018736C56C3][32][192.168.1.89]/24
                      ::                                    32768 ?
 *>i  [2][10.2.2.20:1][0][48][0059DC50AE01][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][0059DC50AE4C][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][0059DC50AE4C][32][192.168.1.81]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][0059DC50AE4C][128][2001:DB8:201::81]/36
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][0059DC50AE4C][128][FE80::259:DCFF:FE50:AE4C]/36
                      10.1.1.10                0    100      0 ?
 *>   [2][10.2.2.20:1][0][48][380E4D9B6A48][32][192.168.1.201]/24
                      ::                                    32768 ?
 *>   [2][10.2.2.20:1][0][48][380E4D9B6A48][128][2001:DB8:201::201]/36
                      ::                                    32768 ?
 *>i  [2][10.2.2.20:1][0][48][A0F84910BCCC][32][192.168.1.201]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:1][0][48][A0F84910BCCC][128][2001:DB8:201::201]/36
                      10.1.1.10                0    100      0 ?
Route Distinguisher: 10.2.2.20:2
 *>   [2][10.2.2.20:2][0][48][0018736C5681][0][*]/20
                      ::                                    32768 ?
 *>   [2][10.2.2.20:2][0][48][0018736C56C2][0][*]/20
                      ::                                    32768 ?
 *>   [2][10.2.2.20:2][0][48][0018736C56C2][32][192.168.2.89]/24
                      ::                                    32768 ?
 *>i  [2][10.2.2.20:2][0][48][0059DC50AE01][0][*]/20
                      10.1.1.10                0    100      0 ?
```

```
 *>i  [2][10.2.2.20:2][0][48][0059DC50AE42][0][*]/20
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:2][0][48][0059DC50AE42][32][192.168.2.81]/24
                      10.1.1.10                0    100      0 ?
 *>   [2][10.2.2.20:2][0][48][380E4D9B6A42][32][192.168.2.202]/24
                      ::                                    32768 ?
 *>   [2][10.2.2.20:2][0][48][380E4D9B6A42][128][2001:DB8:202::202]/36
                      ::                                    32768 ?
 *>i  [2][10.2.2.20:2][0][48][A0F84910BCC2][32][192.168.2.202]/24
                      10.1.1.10                0    100      0 ?
 *>i  [2][10.2.2.20:2][0][48][A0F84910BCC2][128][2001:DB8:202::202]/36
                      10.1.1.10                0    100      0 ?
Route Distinguisher: 103:2
 *>i  [5][103:2][0][24][10.1.13.0]/17
                      10.1.1.10                0    100      0 ?
 *>i  [5][103:2][0][24][192.168.1.0]/17
                      10.1.1.10                0    100      0 ?
 *>i  [5][103:2][0][24][192.168.2.0]/17
                      10.1.1.10                0    100      0 ?
 *>i  [5][103:2][0][64][2001:DB8:201::]/29
                      10.1.1.10                0    100      0 ?
 *>i  [5][103:2][0][64][2001:DB8:202::]/29
                      10.1.1.10                0    100      0 ?
Route Distinguisher: 104:2 (default for vrf green)
 *>   [5][104:2][0][24][10.1.14.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][104:2][0][24][192.168.1.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][104:2][0][24][192.168.2.0]/17
                      0.0.0.0                  0         32768 ?
 *>   [5][104:2][0][64][2001:DB8:201::]/29
                      ::                       0         32768 ?
 *>   [5][104:2][0][64][2001:DB8:202::]/29
                      ::                       0         32768 ?
```

show ip route vrf green

VTEP 1

The following example shows the output for the **show ip route vrf** *vrf-name* command on VTEP 1:

```
VTEP1# show ip route vrf green
Routing Table: green
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, 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, 3 subnets, 2 masks
C        10.1.13.0/24 is directly connected, Loopback13
L        10.1.13.13/32 is directly connected, Loopback13
B        10.1.14.0/24 [200/0] via 10.2.2.20, 01:30:02, Vlan200
      192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks
```

```
C        192.168.1.0/24 is directly connected, Vlan201
B        192.168.1.89/32 [200/0] via 10.2.2.20, 00:04:05, Vlan200
L        192.168.1.201/32 is directly connected, Vlan201
      192.168.2.0/24 is variably subnetted, 3 subnets, 2 masks
C        192.168.2.0/24 is directly connected, Vlan202
B        192.168.2.89/32 [200/0] via 10.2.2.20, 00:04:10, Vlan200
L        192.168.2.202/32 is directly connected, Vlan202
```

VTEP 2

The following example shows the output for the **show ip route vrf** *vrf-name* command on VTEP 2:

```
VTEP2# show ip route vrf green
Routing Table: green
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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, 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, 3 subnets, 2 masks
B        10.1.13.0/24 [200/0] via 10.1.1.10, 01:31:17, Vlan200
C        10.1.14.0/24 is directly connected, Loopback14
L        10.1.14.14/32 is directly connected, Loopback14
      192.168.1.0/24 is variably subnetted, 3 subnets, 2 masks
C        192.168.1.0/24 is directly connected, Vlan201
B        192.168.1.81/32 [200/0] via 10.1.1.10, 01:39:53, Vlan200
L        192.168.1.201/32 is directly connected, Vlan201
      192.168.2.0/24 is variably subnetted, 3 subnets, 2 masks
C        192.168.2.0/24 is directly connected, Vlan202
B        192.168.2.81/32 [200/0] via 10.1.1.10, 01:39:30, Vlan200
L        192.168.2.202/32 is directly connected, Vlan202
```

show platform software fed switch active matm mactable vlan

VTEP 1

The following examples show the output for the **show platform software fed switch active matm mactable vlan** *vlan-id* command on VTEP 1:



Note The MAC address of the peer's core SVI interface must be present in the core VLAN.

```
VTEP1# show platform software fed switch active matm macTable vlan 200
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle           diHandle             *a_time  *e_time  ports

-------------------------------------------------------------------------------------------------------------------------
200    a0f8.4910.bce2      0x8002      0  19880     64  0x7f5d8503fd48      0x7f5d852b6d28
```

```
      0x0                0x5234                      0        0  Vlan200

200    380e.4d9b.6a4a   0x1000001      0      0     64  0x7f5d855bfaa8      0x7f5d852aca68
      0x7f5d851c7078     0x0                         0        0  RLOC 10.2.2.20 adj_id 126


Total Mac number of addresses:: 2
```

```
VTEP1# show platform software fed switch active matm macTable vlan 201
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle           siHandle
      riHandle            diHandle             *a_time  *e_time  ports
------------------------------------------------------------------------------------------------------------------------
201    00aa.00bb.00cc      0x8002      0  42949     64  0x7f5d85007b88      0x7f5d852b6d28
      0x0                0x0                          0        0  Vlan201

201    0059.dc50.ae01         0x1      9      0      0  0x7f5d852abaf8      0x7f5d85035248
      0x0                0x7f5d8517eae8             300        9  GigabitEthernet1/0/1

201    a0f8.4910.bccc      0x8002      0  19880     64  0x7f5d852ad618      0x7f5d852b6d28
      0x0                0x5234                       0        9  Vlan201

201    0059.dc50.ae4c         0x1     16      0      0  0x7f5d855b3ff8      0x7f5d855a2858
      0x0                0x7f5d8517eae8             300       95  GigabitEthernet1/0/1

201    380e.4d9b.6a48      0x8002      0      0     64  0x7f5d84fbf948      0x7f5d852b6d28
      0x0                0x5234                       0       95  Vlan201

201    0018.736c.56c3   0x1000001      0      0     64  0x7f5d855c8268      0x7f5d852368b8
      0x7f5d855c6098     0x0                          0       95  RLOC 10.2.2.20 adj_id 36


Total Mac number of addresses:: 6
```

```
VTEP1# show platform software fed switch active matm macTable vlan 202
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle           siHandle
      riHandle            diHandle             *a_time  *e_time  ports
------------------------------------------------------------------------------------------------------------------------
202    a0f8.4910.bcc2      0x8002      0  19880     64  0x7f5d8503d288      0x7f5d852b6d28
      0x0                0x0                          0        0  Vlan202

202    0059.dc50.ae01         0x1     10      0      0  0x7f5d852ac8b8      0x7f5d852ac668
      0x0                0x7f5d8517eae8             300       15  GigabitEthernet1/0/1

202    0018.736c.5681   0x1000001      0      0     64  0x7f5d855ba7a8      0x7f5d855b0c58
      0x7f5d8518dea8     0x0                          0       15  RLOC 10.2.2.20 adj_id 125

202    0059.dc50.ae42         0x1     17      0      0  0x7f5d8518e848      0x7f5d855a5258
      0x0                0x7f5d8517eae8             300      225  GigabitEthernet1/0/1

202    380e.4d9b.6a42      0x8002      0      0     64  0x7f5d855a59a8      0x7f5d852b6d28
      0x0                0x5234                       0      225  Vlan202

202    0018.736c.56c2   0x1000001      0      0     64  0x7f5d8523d2b8      0x7f5d855b0c58
      0x7f5d8518dea8     0x0                          0      225  RLOC 10.2.2.20 adj_id 125


Total Mac number of addresses:: 6
```

VTEP 2

The following examples show the output for the **show platform software fed switch active matm mactable vlan** *vlan-id* command on VTEP 2:



Note The MAC address of the peer's core SVI interface must be present in the core VLAN.

```
VTEP2# show platform software fed switch active matm macTable vlan 200
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle           diHandle             *a_time  *e_time  ports

-------------------------------------------------------------------------------------------------------------------
200    380e.4d9b.6a4a      0x8002      0    128     64  0x7fa88557f3a8      0x7fa885574e38
      0x0                0x5174                      0        0  Vlan200

200    a0f8.4910.bce2   0x1000001      0      0     64  0x7fa8859a3d38      0x7fa885947ba8
      0x7fa88598bfb8      0x0                       0        0  RLOC 10.1.1.10 adj_id 155


Total Mac number of addresses:: 2


VTEP2# show platform software fed switch active matm macTable vlan 201
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle           diHandle             *a_time  *e_time  ports

-------------------------------------------------------------------------------------------------------------------
201    380e.4d9b.6a48      0x8002      0  42949     64  0x7fa885970018      0x7fa885574e38
      0x0                0x5174                      0        0  Vlan201

201    0059.dc50.ae01   0x1000001      0      0     64  0x7fa8849e1be8      0x7fa88598da48
      0x7fa88598e1f8      0x0                       0        0  RLOC 10.1.1.10 adj_id 153

201    0059.dc50.ae4c   0x1000001      0      0     64  0x7fa885993e68      0x7fa88598da48
      0x7fa88598e1f8      0x0                       0        0  RLOC 10.1.1.10 adj_id 153

201    a0f8.4910.bccc      0x8002      0      0     64  0x7fa8859acc48      0x7fa885574e38
      0x0                0x5174                      0        0  Vlan201

201    0018.736c.56c3         0x1     68      0      0  0x7fa8859d3908      0x7fa88599e108
      0x0                0x7fa884f079d8            300      247  GigabitEthernet1/0/1


Total Mac number of addresses:: 5


VTEP2# show platform software fed switch active matm macTable vlan 202
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle           diHandle             *a_time  *e_time  ports

-------------------------------------------------------------------------------------------------------------------
202    380e.4d9b.6a42      0x8002      0  19018     64  0x7fa885994cd8      0x7fa885574e38
      0x0                0x5174                      0        0  Vlan202

202    0018.736c.5681         0x1      9      0      0  0x7fa88599c4e8      0x7fa88599c218
      0x0                0x7fa884f079d8            300        7  GigabitEthernet1/0/1

202    0059.dc50.ae01   0x1000001      0      0     64  0x7fa8859a3098      0x7fa8859a2dc8
      0x7fa88599ee48      0x0                       0        7  RLOC 10.1.1.10 adj_id 154
```

```
202     0059.dc50.ae42    0x1000001       0       0      64  0x7fa8849e6b78      0x7fa8859a2dc8
      0x7fa88599ee48      0x0                              0        7  RLOC 10.1.1.10 adj_id 154

202     a0f8.4910.bcc2      0x8002        0       0      64  0x7fa88594ddb8      0x7fa885574e38
      0x0                 0x5174                           0        7  Vlan202

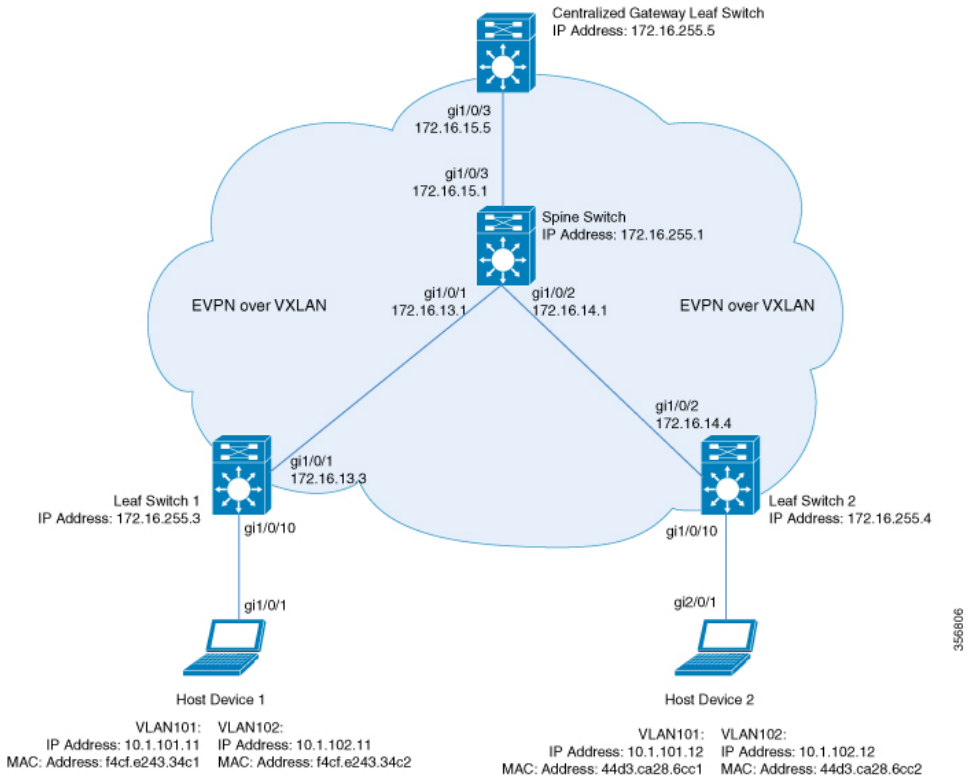
202     0018.736c.56c2         0x1     67       0       0  0x7fa8859d3488      0x7fa8859834f8
      0x0                 0x7fa884f079d8                  300      267  GigabitEthernet1/0/1


Total Mac number of addresses:: 6
```

Example: Enabling EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway

This section provides an example to show how EVPN VXLAN IRB is configured using centralized default gateway. The example shows how to configure and verify EVPN VXLAN IRB for the topology shown in the following image:

Figure 16: EVPN VXLAN Topology with Centralized Default Gateway



The topology shows an EVPN VXLAN network with a spine switch and three leaf switches (Leaf Switch 1, Leaf Switch 2, and Centralized Gateway Leaf Switch). The Centralized Gateway Leaf Switch provides the default gateway for all the VTEPs in the network. EVPN routes of route type 2 are advertised using the Default Gateway Extended Community attribute.

Note

- Each host device in the image is shown with two different VLANs only to illustrate the example.
- The Centralized Gateway Leaf Switch can be configured either as a single switch or a stack of switches.

The following tables provide sample configurations for the devices in the topology above:

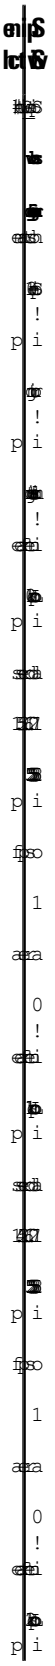
Table 12: Configuring the Leaf Switches and Centralized Gateway Leaf Switch to Enable IRB using Centralized Default Gateway

Leaf Switch 1	Centralized Gateway Switch	Leaf Switch 2
Leaf_Switch1# **show running-config** hostname Leaf_Switch1 ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type static ! l2vpn evpn instance 102 vlan-based encapsulation vxlan replication-type ingress ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport mode trunk ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.3	CGW# **show running-config** hostname CGW ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type static ! l2vpn evpn instance 102 vlan-based encapsulation vxlan replication-type ingress ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.5 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.5 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 !	Leaf_Switch2# **show running-config** hostname Leaf_Switch2 ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ! l2vpn evpn replication-type static router-id Loopback1 ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan replication-type ingress ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport mode trunk ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.4

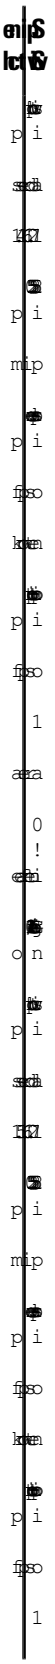
Leaf Switch 1	Centralized Gateway Switch	Leaf Switch 2
!	interface GigabitEthernet1/0/3	!
router bgp 65001	no switchport	router bgp 65001
bgp log-neighbor-changes	ip address 172.16.15.5 255.255.255.0	bgp log-neighbor-changes
no bgp default ipv4-unicast	ip pim sparse-mode	no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001	ip ospf network point-to-point	neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source	ip ospf 1 area 0	neighbor 172.16.255.1 update-source
Loopback0	!	Loopback0
!	interface GigabitEthernet1/0/1	!
address-family ipv4	switchport mode trunk	address-family ipv4
exit-address-family	!	exit-address-family
!	interface Vlan101	!
address-family l2vpn evpn	description centralized gateway	address-family l2vpn evpn
neighbor 172.16.255.1 activate	vrf forwarding green	neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community	ip address 10.1.101.1 255.255.255.0	neighbor 172.16.255.1 send-community
both	!	both
exit-address-family	interface Vlan102	exit-address-family
!	description centralized gateway	!
ip pim rp-address 172.16.255.255	vrf forwarding green	ip pim rp-address 172.16.255.255
!	ip address 10.1.102.1 255.255.255.0	!
end	!	end
!	interface nve1	!
Leaf_Switch1#	no ip address	Leaf_Switch2#
	source-interface Loopback1	
	host-reachability protocol bgp	
	member vni 10101 mcast-group	
	225.0.0.101	
	member vni 10102 ingress-replication	
	!	
	router ospf 1	
	router-id 172.16.255.5	
	!	
	router bgp 65001	
	bgp log-neighbor-changes	
	no bgp default ipv4-unicast	
	neighbor 172.16.255.1 remote-as 65001	
	neighbor 172.16.255.1 update-source	
	Loopback0	
	!	
	address-family ipv4	
	redistribute connected	
	redistribute static	
	exit-address-family	
	!	
	address-family l2vpn evpn	
	neighbor 172.16.255.1 activate	
	neighbor 172.16.255.1 send-community	
	both	
	exit-address-family	
	!	
	address-family ipv4 vrf green	
	advertise l2vpn evpn	
	redistribute connected	
	redistribute static	
	exit-address-family	
	!	
	ip pim rp-address 172.16.255.255	
	!	
	end	

Table 13: Configuring the Spine Switch to Enable IRB using Centralized Default Gateway



















The following examples provide sample outputs for **show** commands on Leaf Switch 1, Leaf Switch 2, and Centralized Gateway Leaf Switch to verify IRB for the topology configured above:

Leaf Switch 1

The following example shows the output for the **show l2vpn evpn default-gateway** command on Leaf Switch 1:

```
Leaf_Switch1# show l2vpn evpn default-gateway
Valid Default Gateway Address   EVI   VLAN  MAC Address    Source
----- ------------------------- ----- ----- -------------- --------------------
  Y   10.1.101.1                101   101   7c21.0dbd.2741 172.16.254.5
  Y   10.1.102.1                102   102   7c21.0dbd.274d 172.16.254.5

Leaf-01#sh l2vpn evpn default-gateway  detail
Default Gateway Address:   10.1.101.1
EVPN Instance:             101
Vlan:                      101
MAC Address:               7c21.0dbd.2741
Ethernet Tag ID:           0
Source:                    V:10101 172.16.254.5

Default Gateway Address:   10.1.102.1
EVPN Instance:             102
Vlan:                      102
MAC Address:               7c21.0dbd.274d
Ethernet Tag ID:           0
Source:                    V:10102 172.16.254.5
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on Leaf Switch 1 for route type 2 and the IP address of the VLAN interface on Leaf Switch 1:

```
Leaf_Switch1# show bgp l2vpn evpn route-type 2 0 7c210dbd2741 10.1.101.1
BGP routing table entry for [2][172.16.254.3:101][0][48][7C210DBD2741][32][10.1.101.1]/24,
 version 31009
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
(global)
    172.16.254.5 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10101
```

```
      Extended Community: RT:65001:101 ENCAP:8 EVPN DEF GW:0:0
      Originator: 172.16.255.5, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F575D9E86B0, path: 0x7F575FBC5B10, pathext: 0x7F575DA095E8, exp_net:
0x7F575D9E8810
      flags: net: 0x0, path: 0x4000000000003, pathext: 0x81
      Updated on Jun 19 2020 12:43:11 UTC
BGP routing table entry for [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24,
 version 31007
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.5 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10101
      Extended Community: RT:65001:101 ENCAP:8 EVPN DEF GW:0:0
      Originator: 172.16.255.5, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F575D9E8810, path: 0x7F575FBC4958, pathext: 0x7F575DA09828
      flags: net: 0x0, path: 0x3, pathext: 0x81
      Updated on Jun 19 2020 12:43:11 UTC
```

The following example shows the output for the **show l2route evpn mac ip** command on Leaf Switch 1:

```
Leaf_Switch1# show l2route evpn mac ip
EVI       ETag  Prod    Mac Address          Host IP                Next Hop(s)
----- ---------- ----- -------------- --------------- --------------------------
  101          0   BGP 44d3.ca28.6cc1     10.1.101.12       V:10101 172.16.254.4
  101          0   BGP 7c21.0dbd.2741      10.1.101.1       V:10101 172.16.254.5
  101          0 L2VPN f4cf.e243.34c1     10.1.101.11               Gi1/0/1:101
  102          0   BGP 44d3.ca28.6cc2     10.1.102.12       V:10102 172.16.254.4
  102          0   BGP 7c21.0dbd.274d      10.1.102.1       V:10102 172.16.254.5
  102          0 L2VPN f4cf.e243.34c2     10.1.102.11               Gi1/0/10:102
```

The following example shows the output for the **show l2fib bridge-domain** command on Leaf Switch 1:

```
Leaf_Switch1# show l2fib bridge-domain 101 address unicast 7c21.0dbd.2741

MAC Address                   : 7c21.0dbd.2741
Reference Count               : 1
Epoch                         : 0
Producer                      : BGP
Flags                         : None
Adjacency                     : VXLAN_UC  PL:2863(1) T:VXLAN_UC [MAC]10101:172.16.254.5
PD Adjacency                  : VXLAN_UC  PL:2863(1) T:VXLAN_UC [MAC]10101:172.16.254.5
```

Leaf Switch 2

The following example shows the output for the **show l2vpn evpn default-gateway** command on Leaf Switch 2:

```
Leaf_Switch2# show l2vpn evpn default-gateway
Valid Default Gateway Address   EVI   VLAN  MAC Address    Source
```

```
----- ------------------------- ----- ----- -------------- --------------------
  Y   10.1.101.1                101   101   7c21.0dbd.2741 172.16.254.5
  Y   10.1.102.1                102   102   7c21.0dbd.274d 172.16.254.5

Leaf-02#sh l2vpn evpn default-gateway  detail
Default Gateway Address:   10.1.101.1
EVPN Instance:             101
Vlan:                      101
MAC Address:               7c21.0dbd.2741
Ethernet Tag ID:           0
Source:                    V:10101 172.16.254.5

Default Gateway Address:   10.1.102.1
EVPN Instance:             102
Vlan:                      102
MAC Address:               7c21.0dbd.274d
Ethernet Tag ID:           0
Source:                    V:10102 172.16.254.5
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on Leaf Switch 2 for route type 2 and the IP address of the VLAN interface on Leaf Switch 1:

```
Leaf_Switch2# show l2vpn evpn route-type 2 0 7c210dbd2741 10.1.101.1
BGP routing table entry for [2][172.16.254.4:101][0][48][7C210DBD2741][32][10.1.101.1]/24,
 version 17202
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
(global)
    172.16.254.5 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10101
      Extended Community: RT:65001:101 ENCAP:8 EVPN DEF GW:0:0
      Originator: 172.16.255.5, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F84B88F13F0, path: 0x7F84BB57B4C8, pathext: 0x7F84B8F48C48, exp_net:
0x7F84B88F0210
      flags: net: 0x0, path: 0x4000000000003, pathext: 0x81
      Updated on Jun 19 2020 12:47:39 UTC
BGP routing table entry for [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24,
 version 17200
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.5 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10101
      Extended Community: RT:65001:101 ENCAP:8 EVPN DEF GW:0:0
      Originator: 172.16.255.5, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F84B88F0210, path: 0x7F84BB57AC58, pathext: 0x7F84B8F48E28
      flags: net: 0x0, path: 0x3, pathext: 0x81
      Updated on Jun 19 2020 12:47:39 UTC
```

The following example shows the output for the **show l2route evpn mac ip** command on Leaf Switch 2:

```
Leaf_Switch2# show l2route evpn mac ip
```

```
  EVI       ETag  Prod    Mac Address         Host IP                Next Hop(s)
----- ---------- ----- -------------- --------------- --------------------------
  101          0 L2VPN 44d3.ca28.6cc1     10.1.101.12               Gi1/0/1:101
  101          0   BGP 7c21.0dbd.2741      10.1.101.1       V:10101 172.16.254.5
  101          0   BGP f4cf.e243.34c1     10.1.101.11       V:10101 172.16.254.3
  102          0 L2VPN 44d3.ca28.6cc2     10.1.102.12               Gi1/0/3:102
  102          0   BGP 7c21.0dbd.274d      10.1.102.1       V:10102 172.16.254.5
  102          0   BGP f4cf.e243.34c2     10.1.102.11       V:10102 172.16.254.3
```

The following example shows the output for the **show l2fib bridge-domain** command on Leaf Switch 2:

```
Leaf_Switch2# show l2fib bridge-domain 101 address unicast 7c21.0dbd.2741

MAC Address                   : 7c21.0dbd.2741
Reference Count               : 1
Epoch                         : 0
Producer                      : BGP
Flags                         : None
Adjacency                     : VXLAN_UC  PL:831(1) T:VXLAN_UC [MAC]10101:172.16.254.5
PD Adjacency                  : VXLAN_UC  PL:831(1) T:VXLAN_UC [MAC]10101:172.16.254.5
```

Centralized Gateway Leaf Switch

The following example shows the output for the **show l2vpn evpn default-gateway** command on Centralized Gateway Leaf Switch:

```
CGW# show l2vpn evpn default-gateway
Valid Default Gateway Address   EVI   VLAN  MAC Address    Source
----- ------------------------ ----- ----- -------------- --------------------
  Y   10.1.101.1               101   101   7c21.0dbd.2741 Vl101
  Y   10.1.102.1               102   102   7c21.0dbd.274d Vl102

CGW#sh l2vpn evpn default-gateway  detail
Default Gateway Address:   10.1.101.1
EVPN Instance:             101
Vlan:                      101
MAC Address:               7c21.0dbd.2741
Ethernet Tag ID:           0
Source:                    V:10101 Vlan101

Default Gateway Address:   10.1.102.1
EVPN Instance:             102
Vlan:                      102
MAC Address:               7c21.0dbd.274d
Ethernet Tag ID:           0
Source:                    V:10102 Vlan102
```

The following example shows the output for the **show bgp l2vpn evpn evi** *evpn-instance* **route-type** command on Centralized Gateway Leaf Switch for route type 2 and the IP address of the VLAN interface on Leaf Switch 1:

```
CGW# show bgp l2vpn evpn evi 101 route-type 2 0 7c210dbd2741 10.1.101.1
BGP routing table entry for [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24,
 version 39
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
```

```
    1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.5)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 10101
      Extended Community: RT:65001:101 ENCAP:8 EVPN DEF GW:0:0
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.5
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F3805208AF0, path: 0x7F380521B380, pathext: 0x7F3806746D98
      flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
      Updated on Jun 19 2020 12:46:25 UTC
```

The following example shows the output for the **show l2route evpn mac ip** command on Centralized Gateway Leaf Switch:

```
CGW# show l2route evpn mac ip
  EVI       ETag  Prod    Mac Address         Host IP                 Next Hop(s)
----- ---------- ----- -------------- --------------- --------------------------
  101          0   BGP 44d3.ca28.6cc1     10.1.101.12       V:10101 172.16.254.4
  101          0 L2VPN 7c21.0dbd.2741      10.1.101.1                    Vl101:0
  101          0   BGP f4cf.e243.34c1     10.1.101.11       V:10101 172.16.254.3
  102          0   BGP 44d3.ca28.6cc2     10.1.102.12       V:10102 172.16.254.4
  102          0 L2VPN 7c21.0dbd.274d      10.1.102.1                    Vl102:0
  102          0   BGP f4cf.e243.34c2     10.1.102.11       V:10102 172.16.254.3
```

The following example shows the output for the **show l2route evpn default-gateway detail** command on Centralized Gateway Leaf Switch:

```
CGW# show l2route evpn default-gateway detail
Ethernet Tag:             0          EVPN Instance:            101

Producer Name:            L2VPN
MAC Address:              7c21.0dbd.2741
Host IP:                  10.1.101.1
Sequence Number:          0
ESI:                      0000.0000.0000.0000.0000
MAC Route Flags:          BInt()Dgl
Next Hop(s):              Vl101:0

EVPN Instance:            102
Ethernet Tag:             0
Producer Name:            L2VPN
MAC Address:              7c21.0dbd.274d
Host IP:                  10.1.102.1
Sequence Number:          0
ESI:                      0000.0000.0000.0000.0000
MAC Route Flags:          BInt()Dgl
Next Hop(s):              Vl102:0
```



CHAPTER 7

Configuring DHCP Relay in a BGP EVPN VXLAN Fabric

- Information About DHCP Relay in a BGP EVPN VXLAN Fabric, on page 155
- How to Configure DHCP Relay in a BGP EVPN VXLAN Fabric, on page 157
- Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 162

Information About DHCP Relay in a BGP EVPN VXLAN Fabric

Networks use DHCP relay to forward DHCP packets between host devices and a DHCP server. In a BGP EVPN VXLAN fabric, you can configure a VTEP as a relay agent to provide DCHP relay services in a multi-tenant VXLAN environment.

When a network uses DHCP relay, DHCP messages move through the same switch in both directions. DHCP relay generally uses the gateway IP address (GiAddr) for scope selection and DHCP response messages. In a BGP EVPN VXLAN fabric that has distributed IP anycast gateway enabled, DHCP messages can return to any switch that hosts the respective GiAddr.

Deploying DHCP relay in an EVPN VXLAN network requires a different method for scope selection and a unique IP address for each switch in the network. The unique Loopback interface for a switch becomes the GiAddr that a switch uses to respond to the correct switch. DHCP option 82, also referred to as DHCP option VPN, is used for scope selection based on the Layer 2 VNI.

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

- **Sub-Option 151(0x97)—Virtual Subnet Selection:**

  The virtual subnet selection sub-option is used to convey VRF-related information to the DHCP server in an MPLS VPN and a VXLAN EVPN multi-tenant environment.

  RFC 6607 provides the definition for this sub-option.

- **Sub-Option 11(0xb)—Server ID Override**

  The server identifier or server ID override sub-option allows the DHCP relay agent to specify a new value for the server ID option. The DHCP server inserts this new value in the reply packet. This sub-option allows the DHCP relay agent to act as the actual DHCP server. The DHCP relay agent begins to receive all the renew requests instead of the DHCP server. The server ID override sub-option contains the incoming interface IP address. The DHCP client accesses the DHCP relay agent using the incoming interface IP address. The DHCP client uses this information to send all the renew and release request

packets to the DHCP relay agent. The DHCP relay agent adds all 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). To implement the suboption and manage the function, run the **ip dhcp relay sub-option type cisco** command in global configuration mode on the VTEP that acts as the DHCP relay agent.

RFC 5107 provides the definition for this sub-option.

- **Sub-Option 5(0x5)—Link Selection:**

  The link selection sub-option provides a mechanism to separate the subnet or link, on which the DHCP client resides, from the GiAddr. The DHCP server uses this mechanism to communicate with the DHCP relay agent. The DHCP relay agent sets the sub-option to the correct subscriber subnet. The DHCP server then uses this value to assign an IP address different from the GiAddr. The DHCP relay agent sets the GiAddr to its own IP address to ensure that it is possible to forward the DHCP messages over the network.

  For this function, Cisco's proprietary implementation is sub-option 150(0x96). To manage the function, run the **ip dhcp relay sub-option type cisco** command in global configuration mode on the VTEP that acts as the DHCP relay agent.

  RFC 3527 provides the definition for this sub-option.

DHCP Relay on VTEPs in Distributed Anycast Gateway Deployment

DHCP relay is generally configured on the default gateway that faces the DHCP client. You can configure a VTEP as a DHCP relay agent in different ways to automate IP addressing. The configuration depends on whether the DHCP server is present in the same network, the same VRF, or a different VRF compared to the DHCP client. When the DHCP server and DHCP client are in different VRFs, traffic is forwarded across the tenant or VRF boundaries.

The following are the common DHCP relay deployment scenarios for a BGP EVPN VXLAN fabric:

1. DHCP server is in the Layer 3 default VRF and DHCP client is in the tenant VRF.

   See Example: DHCP Server is in the Layer 3 Default VRF and the DHCP Client is in the Tenant VRF, on page 163 for a configuration example.

2. DHCP server and DHCP client are in the same tenant VRF.

   See Example: DHCP Server and DHCP Client are in the Same Tenant VRF, on page 167 for a configuration example.

3. DHCP server and DHCP client are in different tenant VRFs.

   See Example: DHCP Client and DHCP Server are in Different Tenant VRFs, on page 170 for a configuration example.

4. DHCP server is in a non-default non-VXLAN VRF and DHCP client is in the tenant VRF.

   See Example: DHCP Server is in a non-Default, non-VXLAN VRF and DHCP Client is in the Tenant VRF, on page 175 for a configuration example.

DHCP Relay on VTEPs in a Layer 2 Overlay Fabric

In an EVPN VXLAN centralized gateway (CGW) deployment, the CGW performs the Layer 3 gateway function for all the Layer 2 VNIs. All the other Leaf VTEPs in the network perform only bridging. The CGW VTEP acts as the Layer 3 gateway and performs routing for the intersubnet VXLAN traffic.

One of the VTEPs, usually a Border VTEP, is configured with an SVI to act as the Layer 3 gateway. You can also configure the Layer 3 gateway on an external router or a firewall that is connected to the Border Layer 2 VTEP, for policy enforcement and inspection of the intersubnet traffic.

DHCP follows the Discover, Offer, Request, Acknowledge (DORA) process to assign the IP addresses to the endpoints or hosts. To optimize the DORA sequence, a Layer 2 VTEP forwards the DHCP requests only to a Layer 3 VTEP. The Layer 2 VTEP uses the Gateway MAC-IP route (Route type 2 that is received with BGP gateway extended community attribute) to forward the packets to the Layer 3 gateway VTEP. The DHCP Relay on the Layer 3 gateway VTEP communicates with the DHCP server.

The following two configurations are supported.

- Layer 3 Gateway configured on the Border VTEP
- Layer 3 Gateway configured on an external router or a firewall

How to Configure DHCP Relay in a BGP EVPN VXLAN Fabric

You must configure EVPN VXLAN Layer 2 and Layer 3 overlay networks before configuring BGP EVPN VXLAN interworking with DHCP relay. See Configuring EVPN VXLAN Integrated Routing and Bridging for detailed steps.

Perform the following set of procedures to configure BGP EVPN VLAN interworking with DHCP relay:

Configuring DHCP Relay on a VTEP

To configure DHCP relay on a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip dhcp relay information option vpn** **Example:** `Device(config)# ip dhcp relay information option vpn`	Adds option VPN suboption to DHCP option 82. Enables the device to insert VPN suboptions into the DHCP relay agent information option

	Command or Action	Purpose
		in the messages forwarded to the DHCP server and sets the GiAddr on the outgoing interface towards the DHCP server.
Step 4	**ip dhcp relay information option** **Example:** `Device(config)# ip dhcp relay information option`	Enables DHCP option 82. Enables the system to insert a DHCP relay agent information option in the messages forwarded to the DHCP server.
Step 5	**ip dhcp relay override** *gateway-ip-address* **link-selection** **Example:** `Device(config)# ip dhcp relay override giaddr link-selection`	Sets the gateway IP address as the IP address of the DHCP relay agent and configures the server to assign an IP address that is different from the GiAddr to the DHCP clients.

Step 6	**ip dhcp compatibility suboption** {**link-selection** | **server-override**} **standard**
	Example:
	`Device(config)# ip dhcp compatibility suboption link-selection standard`
	`Device(config)# ip dhcp compatibility suboption server-override standard`
		Configures the DHCP client to use the Internet Assigned Numbers Authority (IANA) standard relay agent server ID override suboption.
		Use the **link-selection standard** keyword to switch to standard DHCP option 82[5].
		Use the **server-override standard** keyword to switch to standard DHCP option 82[11].

Step 7	**ip dhcp snooping vlan** *vlan-id-list* **Example:** `Device(config)# ip dhcp snooping vlan 201-202`	Enables DHCP snooping on the specified list of VLANs. **Note** VLAN id used for core SVI (L3 VNI) should not be enabled with DHCP snooping.
Step 8	**ip dhcp snooping** **Example:** `Device(config)# ip dhcp snooping`	Enables DHCP snooping globally on the VTEP.
Step 9	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuring DHCP Relay on the Access SVI of a VTEP

Perform this procedure on all the VTEPs for each VLAN that is associated with the Layer 2 VNI configured in the EVPN VXLAN network.

To configure DHCP relay on the access SVI of a VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *vlan-id* **Example:** `Device(config)# interface Vlan 201`	Enters interface configuration mode for the specified VLAN interface. This VLAN interface acts as the GiAddr.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Associates the VRF with the interface. The interface must be associated with the same VRF for which the Layer 3 VNI has been configured for the EVPN VXLAN network.
Step 5	**ip dhcp relay information option vpn-id** **Example:** `Device(config-if)# ip dhcp relay information option vpn-id`	Enables the device to insert VPN suboptions into the DHCP relay agent information option in the messages forwarded to the DHCP server and sets the GiAddr on the outgoing interface towards the DHCP server.
Step 6	**ip dhcp relay source-interface Loopback** *loopback-interface-id* **Example:** `Device(config-if)# ip dhcp relay source-interface Loopback13`	Configures the specified Loopback interface as the source interface for DHCP relay messages. The DHCP relay agent uses the IP address of the source interface as the source IP address to relay messages. **Note** The IP address configured on the Loopback interface must be unique per VTEP per VRF.
Step 7	**ip address** *ip-address* **Example:** `Device(config-if)# ip address 192.168.1.201 255.255.255.0`	Sets the IP address for the VLAN interface.
Step 8	**ip helper-address** [**global** | **vrf** *vrf-name*] *ip-address* **Example:** `Device(config-if)# ip helper-address global 192.168.3.100`	Sets the DHCP IP helper address for the VLAN interface. Use the **global** keyword if the DHCP server is reachable over the global routing table (GRT). Use the **vrf** *vrf-name* keyword if the DHCP server is reachible over the tenant VRF.

	Command or Action	Purpose
	`Device(config-if)# ip helper-address vrf green 192.168.20.20`	
Step 9	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 10	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuring the Layer 3 or Routed Interface on the Border VTEP for DHCP Server Reachability

DHCP server reachability can be achieved through a physical Layer 3 interface (or subinterface), a dot1Q interface, an SVI, or a Layer 3 Portchannel interface (or subinterface).

Note This task is optional if you implement plain IP address forwarding in the respective VRF.

To configure the Layer 3 or routed interface on the border VTEP for external connectivity, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *vlan-id* **Example:** `Device(config)# interface vlan 203`	Enters interface configuration mode for the specified VLAN interface.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Configures the SVI for the VLAN and associates the specified VRF with the interface.

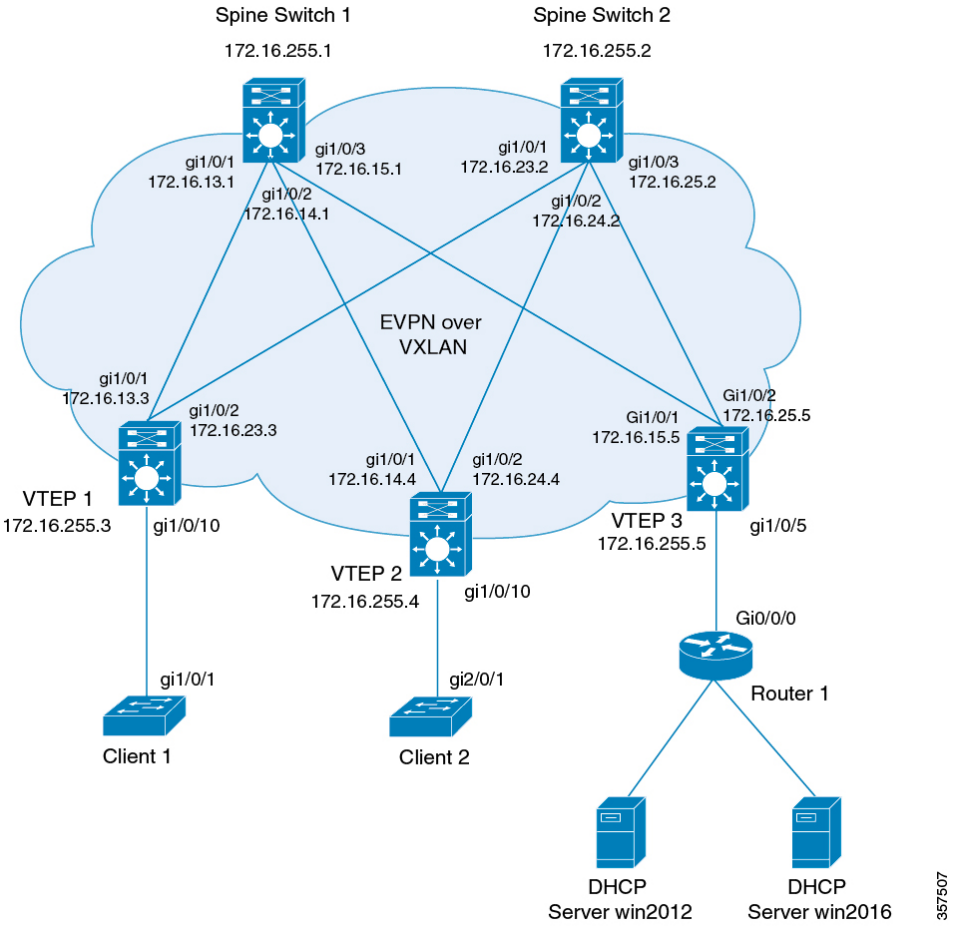
	Command or Action	Purpose
Step 5	**ip address** *ip-address* **Example:** `Device(config-if)# ip address 192.168.3.203 255.255.255.0`	Configures the IP address for the VLAN.
Step 6	**ipv6 address** *ipv6-address* **Example:** `Device(config-if)# ipv6 address 2001:203::203/64`	Configures the IPv6 address for the VLAN.
Step 7	**ipv6 enable** **Example:** `Device(config-if)# ipv6 enable`	Enables IPv6 processing on the VLAN interface.
Step 8	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 9	**interface** *interface-id* **Example:** `Device(config)# interface GigabitEthernet1/0/1`	Enters interface configuration mode for the specified interface.
Step 10	**switchport access vlan** *vlan-id* **Example:** `Device(config-if)# switchport access vlan 203`	Specifies the VLAN to be used as access VLAN when the interface is in access mode.
Step 11	**switchport mode access** **Example:** `Device(config-if)# switchport mode access`	Configures the interface as an access interface.
Step 12	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 13	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric

This section provides configuration examples for DHCP relay in a BGP EVPN VXLAN fabric for the following scenarios using the topology in .

- DHCP server is in the Layer 3 default VRF and DHCP client is in the tenant VRF
- DHCP server and DHCP client are in the same tenant VRF
- DHCP server is in a different tenant VRF from that of the DHCP client
- DHCP server is in a non-default, non-VXLAN VRF and DHCP client is in the tenant VRF

Figure 17: DHCP Relay Deployment in a BGP EVPN VXLAN Fabric



The preceding figure shows an EVPN VXLAN network with two spine switches (Spine Switch 1 and Spine Switch 2) and three leaf switches (VTEP1, VTEP 2, and VTEP 3). VTEP 3 is connected to two DHCP servers. VTEP 1 and VTEP 2 are connected to a single DHCP client each.

Example: DHCP Server is in the Layer 3 Default VRF and the DHCP Client is in the Tenant VRF

This example shows how to configure DHCP relay deployment in a BGP EVPN VXLAN fabric for the topology in this figure when the DHCP server and the DHCP client are in the same tenant VRF. The DHCP server is reachable over global routing table (GRT).

The following tables provide sample configurations for the DHCP server and VTEP 1:

Note VLAN id used for core SVI (L3 VNI) should not be enabled with DHCP snooping.

Table 14: Configuring DHCP when DHCP Server is in the Layer 3 Default VRF and DHCP Client is in the Tenant VRF

DHCP Configuration Snippet

<snip: only the relevant configuration is shown>
 	
ip dhcp-relay source-interface Loopback0
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
ip dhcp snooping vlan 101-102,201-202
ip dhcp snooping
!
interface Loopback0
  ip address 172.16.255.3 255.255.255.255
  ip ospf 1 area 0
!
interface Vlan101
  vrf forwarding green
  ip address 10.1.101.1 255.255.255.0
  ip helper-address global 192.168.20.20
!
interface Vlan102
  vrf forwarding green
  ip address 10.1.102.1 255.255.255.0
  ip helper-address global 192.168.20.20
!
interface Vlan201
  vrf forwarding red
  ip address 10.2.201.1 255.255.255.0
  ip helper-address global 192.168.20.20
!
<snip: only the relevant configuration is shown>

Table 15: Configuring VTEP 1 when DHCP Server is in the Layer 3 Default VRF and DHCP Client is in the Tenant VRF

VTEP 1

```
Leaf-01# show running-config
!
hostname Leaf-01
!
vrf definition green
rd 1:1
!
address-family ipv4
route-target export 1:1
route-target import 1:1
route-target export 1:1 stitching
route-target import 1:1 stitching
exit-address-family
!
vrf definition red
rd 2:2
!
address-family ipv4
route-target export 2:2
route-target import 2:2
route-target export 2:2 stitching
route-target import 2:2 stitching
exit-address-family
!
ip routing
!
ip multicast-routing
!
ip dhcp-relay source-interface Loopback0
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
!
ip dhcp snooping vlan 101-102,201-202
ip dhcp snooping
!
l2vpn evpn
replication-type static
router-id Loopback1
default-gateway advertise
!
l2vpn evpn instance 101 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 102 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 201 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 202 vlan-based
encapsulation vxlan
!
system mtu 9198
!
```

VTEP 1

```
vlan configuration 101
member evpn-instance 101 vni 10101
vlan configuration 102
member evpn-instance 102 vni 10102
vlan configuration 201
member evpn-instance 201 vni 10201
vlan configuration 202
member evpn-instance 202 vni 10202
vlan configuration 901
member vni 50901
vlan configuration 902
member vni 50902
!
interface Loopback0
ip address 172.16.255.3 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.3 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0

!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.23.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/10
switchport mode trunk
!
interface Vlan101
vrf forwarding green
ip address 10.1.101.1 255.255.255.0
ip helper-address global 192.168.20.20
!
interface Vlan102
vrf forwarding green
ip address 10.1.102.1 255.255.255.0
ip helper-address global 192.168.20.20
!
interface Vlan201
vrf forwarding red
ip address 10.2.201.1 255.255.255.0
ip helper-address global 192.168.20.20
!
interface Vlan202
vrf forwarding red
ip address 10.2.202.1 255.255.255.0
ip helper-address global 192.168.20.20
!
```

VTEP 1

```
interface Vlan901
vrf forwarding green
ip unnumbered Loopback0
no autostate
!
interface Vlan902
vrf forwarding red
ip unnumbered Loopback0
no autostate
!
!
interface nve1
no ip address
source-interface Loopback1
host-reachability protocol bgp
member vni 10101 mcast-group 225.0.0.101
member vni 10102 mcast-group 225.0.0.102
member vni 10201 mcast-group 225.0.0.201
member vni 10202 mcast-group 225.0.0.202
member vni 50901 vrf green
member vni 50902 vrf red
!
router ospf 1
router-id 172.16.255.3
!
router bgp 65001
bgp router-id interface Loopback0
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.2 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0
!
address-family ipv4
exit-address-family
!
address-family l2vpn evpn
neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate
neighbor 172.16.255.2 send-community both
exit-address-family
!
address-family ipv4 vrf green
advertise l2vpn evpn
redistribute connected
redistribute static
exit-address-family
!
address-family ipv4 vrf red
advertise l2vpn evpn
redistribute connected
redistribute static
exit-address-family
!
ip pim rp-address 172.16.255.255
!
end

Leaf-01#
```

Return to Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 162.

Example: DHCP Server and DHCP Client are in the Same Tenant VRF

This example shows how to configure DHCP relay deployment in a BGP EVPN VXLAN fabric for the topology in this figure when the DHCP server and the DHCP client are in the same tenant VRF. The DHCP server is reachable over this common tenant VRF.

The following tables provide sample configurations for the DHCP server and VTEP 1:

Note VLAN id used for core SVI (L3 VNI) should not be enabled with DHCP snooping.

Table 16: Configuring DHCP when DHCP Server and DHCP Client are in the Same Tenant VRF

DHCP Configuration Snippet

<snip: only the relevant configuration is shown>
 	
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
ip dhcp snooping vlan 101-102,201-202
ip dhcp snooping
!
interface Loopback101
  vrf forwarding green
  ip address 10.1.251.1 255.255.255.255
!
interface Vlan101
  vrf forwarding green
  ip dhcp relay source-interface Loopback101
  ip address 10.1.101.1 255.255.255.0
  ip helper-address 192.168.20.20
!
interface Vlan102
  vrf forwarding green
  ip dhcp relay source-interface Loopback101
  ip address 10.1.102.1 255.255.255.0
  ip helper-address 192.168.20.20
 	
<snip: only the relevant configuration is shown>

Table 17: Configuring VTEP 1 when DHCP Server and DHCP Client are in the Same Tenant VRF

VTEP 1

```
Leaf-01# show running-config
!
hostname Leaf-01
!
vrf definition green
rd 1:1
!
address-family ipv4
route-target export 1:1
route-target import 1:1
route-target export 1:1 stitching
route-target import 1:1 stitching
exit-address-family
!
ip routing
!
ip multicast-routing
!
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
!
ip dhcp snooping vlan 101-102,201-202
ip dhcp snooping
!
l2vpn evpn
replication-type static
router-id Loopback1
default-gateway advertise
!
l2vpn evpn instance 101 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 102 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 201 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 202 vlan-based
encapsulation vxlan
!
system mtu 9198
!
vlan configuration 101
member evpn-instance 101 vni 10101
vlan configuration 102
member evpn-instance 102 vni 10102
vlan configuration 201
member evpn-instance 201 vni 10201
vlan configuration 202
member evpn-instance 202 vni 10202
vlan configuration 901
member vni 50901
vlan configuration 902
member vni 50902
!
```

VTEP 1

```
interface Loopback0
ip address 172.16.255.3 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.3 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface Loopback101
vrf forwarding green
ip address 10.1.251.1 255.255.255.255
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.23.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/10
switchport mode trunk
!
interface Vlan101
vrf forwarding green
ip dhcp relay source-interface Loopback101
ip address 10.1.101.1 255.255.255.0
ip helper-address 192.168.20.20
!
interface Vlan102
vrf forwarding green
ip dhcp relay source-interface Loopback101
ip address 10.1.102.1 255.255.255.0
ip helper-address 192.168.20.20
!
interface Vlan901
vrf forwarding green
ip unnumbered Loopback0
no autostate
!
```

VTEP 1

```
interface nve1
no ip address
source-interface Loopback1
host-reachability protocol bgp
member vni 10101 mcast-group 225.0.0.101
member vni 10102 mcast-group 225.0.0.102
member vni 50901 vrf green
!
router ospf 1
router-id 172.16.255.3
!
router bgp 65001
bgp router-id interface Loopback0
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.2 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0
!
address-family ipv4
exit-address-family
!
address-family l2vpn evpn
neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate
neighbor 172.16.255.2 send-community both
exit-address-family
!
address-family ipv4 vrf green
advertise l2vpn evpn
redistribute connected
redistribute static
exit-address-family
!
ip pim rp-address 172.16.255.255
!
end
!
Leaf-01#
```

Return to Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 162.

Example: DHCP Client and DHCP Server are in Different Tenant VRFs

This example shows how to configure DHCP relay deployment in a BGP EVPN VXLAN fabric for the topology in this figure when the DHCP server and the DHCP client are in different tenant VRFs. The DHCP server is reachable over a VRF that is different from the client's VRF.

The following tables provide sample configurations for the DHCP server and VTEP 1:

Note VLAN id used for core SVI (L3 VNI) should not be enabled with DHCP snooping.

Table 18: Configuring DHCP when DHCP Server and DHCP Client are in Different Tenant VRFs

DHCP Configuration Snippet
<snip: only the relevant configuration is shown> ip dhcp relay information option vpn ip dhcp relay information option ip dhcp compatibility suboption link-selection standard ip dhcp compatibility suboption server-override standard ip dhcp snooping vlan 101-102,201-202 ip dhcp snooping ! interface Loopback101 vrf forwarding green ip address 10.1.251.1 255.255.255.255 ! interface Vlan201 vrf forwarding red ip dhcp relay source-interface Loopback101 ip address 10.2.201.1 255.255.255.0 ip helper-address vrf green 192.168.20.20 **<snip: only the relevant configuration is shown>**

Table 19: Configuring VTEP 1 when DHCP Server and DHCP Client are in Different Tenant VRFs

VTEP 1

```
Leaf-01# show running-config
!
hostname Leaf-01
!
vrf definition green
rd 1:1
!
address-family ipv4
route-target export 1:1
route-target import 1:1
route-target export 1:1 stitching
route-target import 1:1 stitching
exit-address-family
!
vrf definition red
rd 2:2
!
address-family ipv4
route-target export 2:2
route-target import 2:2
route-target export 2:2 stitching
route-target import 2:2 stitching
exit-address-family
!
ip routing
!
ip multicast-routing
!
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
!
ip dhcp snooping vlan 101-102,201-202
ip dhcp snooping
!
l2vpn evpn
replication-type static
router-id Loopback1
default-gateway advertise
!
l2vpn evpn instance 101 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 102 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 201 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 202 vlan-based
encapsulation vxlan
!
system mtu 9198
!
```

VTEP 1

```
vlan configuration 101
member evpn-instance 101 vni 10101
vlan configuration 102
member evpn-instance 102 vni 10102
vlan configuration 201
member evpn-instance 201 vni 10201
vlan configuration 202
member evpn-instance 202 vni 10202
vlan configuration 901
member vni 50901
vlan configuration 902
member vni 50902
!
interface Loopback0
ip address 172.16.255.3 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.3 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0

!
interface Loopback101
vrf forwarding green
ip address 10.1.251.1 255.255.255.255

!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.23.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/10
switchport mode trunk
!

interface Vlan101
vrf forwarding green
ip dhcp relay source-interface Loopback101
ip address 10.1.101.1 255.255.255.0
ip helper-address 192.168.20.20
!
interface Vlan102
vrf forwarding green
ip dhcp relay source-interface Loopback101
ip address 10.1.102.1 255.255.255.0
ip helper-address 192.168.20.20
```

VTEP 1

```
interface Vlan201
vrf forwarding red
ip dhcp relay source-interface Loopback101
ip address 10.2.201.1 255.255.255.0
ip helper-address vrf green 192.168.20.20
!
interface Vlan202
vrf forwarding red
ip dhcp relay source-interface Loopback101
ip address 10.2.202.1 255.255.255.0
ip helper-address vrf green 192.168.20.20
!
interface Vlan901
vrf forwarding green
ip unnumbered Loopback0
no autostate
!
interface Vlan902
vrf forwarding red
ip unnumbered Loopback0
no autostate
!
!
interface nve1
no ip address
source-interface Loopback1
host-reachability protocol bgp
member vni 10101 mcast-group 225.0.0.101
member vni 10102 mcast-group 225.0.0.102
member vni 10201 mcast-group 225.0.0.201
member vni 10202 mcast-group 225.0.0.202
member vni 50901 vrf green
member vni 50902 vrf red
!
router ospf 1
router-id 172.16.255.3
!
router bgp 65001
bgp router-id interface Loopback0
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.2 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0
!
address-family ipv4
exit-address-family
!
address-family l2vpn evpn
neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate
neighbor 172.16.255.2 send-community both
exit-address-family
!
```

VTEP 1

```
address-family ipv4 vrf green
advertise l2vpn evpn
redistribute connected
redistribute static
exit-address-family
!
address-family ipv4 vrf red
advertise l2vpn evpn
redistribute connected
redistribute static
exit-address-family
!
ip pim rp-address 172.16.255.255
!
end
Leaf-01#
```

Return to Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 162.

Example: DHCP Server is in a non-Default, non-VXLAN VRF and DHCP Client is in the Tenant VRF

This example shows how to configure DHCP relay deployment in a BGP EVPN VXLAN fabric for the topology in this figure when the DHCP server is in a non-default, non-VXLAN VRF and the DHCP client is in the tenant VRF. The DHCP server is reachable over a VRF that is different from the client's VRF.

The following tables provide sample configurations for the DHCP server and VTEP 1:

Note VLAN id used for core SVI (L3 VNI) should not be enabled with DHCP snooping.

Table 20: Configuring DHCP when DHCP Server is in a non-Default, non-VXLAN VRF and DHCP Client is in the Tenant VRF

DHCP Configuration Snippet
<snip: only the relevant configuration is shown> ip dhcp relay information option vpn ip dhcp relay information option ip dhcp compatibility suboption link-selection standard ip dhcp compatibility suboption server-override standard ip dhcp snooping vlan 101-102,201-202 ip dhcp snooping ! interface Loopback101 vrf forwarding green ip address 10.1.251.1 255.255.255.255 ! interface Vlan201 vrf forwarding red ip dhcp relay source-interface Loopback101 ip address 10.2.201.1 255.255.255.0 ip helper-address vrf green 192.168.20.20 **<snip: only the relevant configuration is shown>**

Table 21: Configuring VTEP 1 when the DHCP Server is in a non-Default, non-VXLAN VRF and the DHCP Client is in the Tenant VRF

VTEP 1

```
Leaf-01# show running-config
!
hostname Leaf-01
!
vrf definition green
rd 1:1
!
address-family ipv4
route-target export 1:1
route-target import 1:1
route-target export 1:1 stitching
route-target import 1:1 stitching
exit-address-family
!
vrf definition red
rd 2:2
!
address-family ipv4
route-target export 2:2
route-target import 2:2
route-target export 2:2 stitching
route-target import 2:2 stitching
exit-address-family
!
ip routing
!
ip multicast-routing
!
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
!
ip dhcp snooping vlan 101-102,201-202
ip dhcp snooping
!
l2vpn evpn
replication-type static
router-id Loopback1
default-gateway advertise
!
l2vpn evpn instance 101 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 102 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 201 vlan-based
encapsulation vxlan
!
l2vpn evpn instance 202 vlan-based
encapsulation vxlan
!
system mtu 9198
!
vlan configuration 101
member evpn-instance 101 vni 10101
vlan configuration 102
member evpn-instance 102 vni 10102
```

VTEP 1

```
vlan configuration 201
member evpn-instance 201 vni 10201
vlan configuration 202
member evpn-instance 202 vni 10202
vlan configuration 901
member vni 50901
vlan configuration 902
member vni 50902
!
interface Loopback0
ip address 172.16.255.3 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.3 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0

!
interface Loopback101
vrf forwarding green
ip address 10.1.251.1 255.255.255.255

!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.23.3 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/10
switchport mode trunk
!

interface Vlan101
vrf forwarding green
ip dhcp relay source-interface Loopback101
ip address 10.1.101.1 255.255.255.0
ip helper-address 192.168.20.20
!
interface Vlan102
vrf forwarding green
ip dhcp relay source-interface Loopback101
ip address 10.1.102.1 255.255.255.0
ip helper-address 192.168.20.20
```

VTEP 1

```
interface Vlan201
vrf forwarding red
ip dhcp relay source-interface Loopback101
ip address 10.2.201.1 255.255.255.0
ip helper-address vrf green 192.168.20.20
!
interface Vlan202
vrf forwarding red
ip dhcp relay source-interface Loopback101
ip address 10.2.202.1 255.255.255.0
ip helper-address vrf green 192.168.20.20
!
interface Vlan901
vrf forwarding green
ip unnumbered Loopback0
no autostate
!
interface Vlan902
vrf forwarding red
ip unnumbered Loopback0
no autostate
!
!
interface nve1
no ip address
source-interface Loopback1
host-reachability protocol bgp
member vni 10101 mcast-group 225.0.0.101
member vni 10102 mcast-group 225.0.0.102
member vni 10201 mcast-group 225.0.0.201
member vni 10202 mcast-group 225.0.0.202
member vni 50901 vrf green
member vni 50902 vrf red
!
router ospf 1
router-id 172.16.255.3
!
router bgp 65001
bgp router-id interface Loopback0
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.2 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0
!
address-family ipv4
exit-address-family
!
address-family l2vpn evpn
neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate
neighbor 172.16.255.2 send-community both
exit-address-family
!
```

VTEP 1
address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv4 vrf red advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-01#

Return to Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 162.

Example: Layer 3 Gateway Configured on Border VTEP

This section shows how to configure a DHCP Relay in a Layer 2 overlay fabric, when the CGW, which is the Layer 3 gateway, is configured on the border VTEP. In this topology <add the topology figure>, the Layer 2 VTEPs relay the DHCP packets to the Border Layer 3 VTEP.

The following tables provide a snippet of the sample configurations for the Border VTEP (Layer 3 Gateway) and Leaf VTEP.

Note VLAN id used for core SVI (L3 VNI) should not be enabled with DHCP snooping.

Table 22: DHCP Configuration on Leaf VTEP

Leaf Layer 2 VTEP Configuration Snippet
<snip: only the relevant configuration is shown> ip dhcp relay information option vpn ip dhcp relay information option ip dhcp compatibility suboption link-selection standard ip dhcp compatibility suboption server-override standard ip dhcp snooping vlan 101-102 ip dhcp snooping !

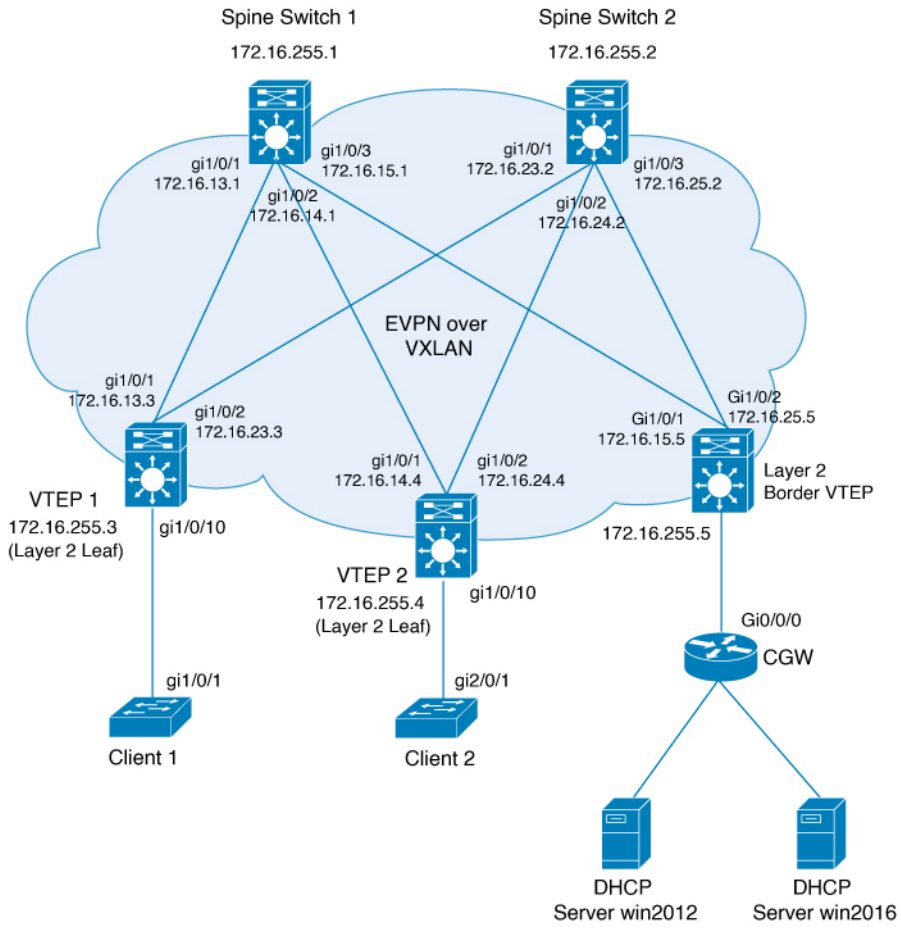
Table 23: DHCP Configuration on Layer 3 Border VTEP

Layer 3 Border VTEP (CGW) Configuration Snippet
<snip: only the relevant configuration is shown> ip dhcp-relay source-interface Loopback0 ip dhcp relay information option vpn ip dhcp relay information option ip dhcp compatibility suboption link-selection standard ip dhcp compatibility suboption server-override standard ip dhcp snooping vlan 101-102,201-202 ip dhcp snooping ! interface Loopback0 ip address 172.16.255.5 255.255.255.255 ip ospf 1 area 0 ! interface Vlan101 description Layer 3 gateway for vni 10101 ip address 10.1.101.1 255.255.255.0 ip helper-address global 192.168.20.20 !

Example: Layer 3 Gateway Configured on External Router or Firewall

This section shows how to configure a DHCP Relay in a Layer 2 overlay fabric, when the CGW, which is the Layer 3 gateway, is configured on an external router or a firewall. In this topology, all the VTEPs perform Layer 2 bridging. A Layer 2 Border VTEP connects to the external router that acts as a Layer 3 gateway. The Layer 2 Border node advertises a Route Type 2 (RT2) with default gateway extended community.

Figure 18: Layer 3 Gateway Configured on External Router



The following tables provide a snippet of the sample configurations for the Layer 2 Border VTEP and Leaf VTEP.

Table 24: DHCP Configuration on Leaf VTEP

Leaf VTEP Configuration Snippet

Leaf VTEP Configuration Snippet

```
<snip: only the relevant configuration is shown>
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
ip dhcp snooping vlan 101-102
ip dhcp snooping

VTEP-2#sh run | s dhcp
ip dhcp snooping vlan 101 #-------------- globally
ip dhcp snooping #----------------------- globally

on DHCP Gateway facing port:
VTEP-2#sh run int GigabitEthernet1/0/1
Building configuration...

Current configuration : 139 bytes
!
interface GigabitEthernet1/0/10
 switchport access vlan 101
 switchport mode access
 spanning-tree portfast
 ip dhcp snooping trust #-----------------make trusted port GW port
end
VTEP-2#sh run | s ip access-list
ip access-list extended gatewayiplist #------------- ACL
 10 permit ip host 192.168.101.1 any
VTEP-2#
VTEP-2#sh run | s route-map
route-map GW_MAC_IP_TRACKING permit 10 #------------Route Map
 match ip address gatewayiplist
 match evpn route-type 2-mac-ip
 set extcommunity default-gw

VTEP-2#sh run | s r b
router bgp 65001
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 neighbor 172.16.255.1 remote-as 65001
 neighbor 172.16.255.1 update-source Loopback0
 !
 address-family ipv4
 exit-address-family
 !
 address-family l2vpn evpn
  neighbor 172.16.255.1 activate
  neighbor 172.16.255.1 send-community both
  neighbor 172.16.255.1 route-map GW_MAC_IP_TRACKING out #----------Apply route Map to
L2VPN neighbor
 exit-address-family


**NOTE: If VRF configured on VTEPs then option 82 is needed, in that case you need below
additional commands on VTEP-2.
If it is "flat" network not needed:
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard

**Genral debugs used**
  debug ip dhcp server packet detail
```

Leaf VTEP Configuration Snippet

```
  debug ip dhcp server packet
  debug ip dhcp server events
  debug ip dhcp snooping packet
  debug dhcp detail

**To verify check on VTEP-1**
VTEP-1#sh l2vpn evpn default-gateway
Valid Default Gateway Address                  EVI   VLAN  MAC Address    Source
----- ---------------------------------------- ----- ----- -------------- -----------
  Y   192.168.101.1                            101   101   4488.16ff.d7c1 172.16.254.2
#---------"Y" is VALID
```

Table 25: DHCP Configuration on Layer 2 Border VTEP

Layer 2 Border VTEP Configuration Snippet

<snip: only the relevant configuration is shown>

```
ip dhcp relay information option vpn
ip dhcp relay information option
ip dhcp compatibility suboption link-selection standard
ip dhcp compatibility suboption server-override standard
ip dhcp snooping vlan 101-102
ip dhcp snooping

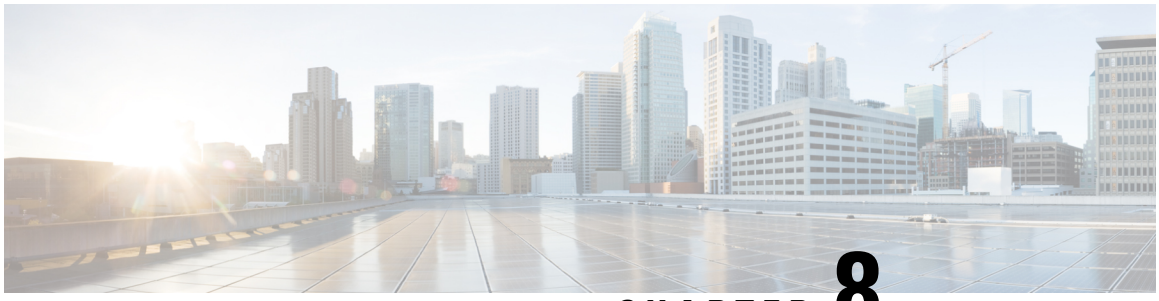
    ! Configure an access list and route-map to match and set default-gw attribute

ip access-list extended gatewayiplist
 10 permit ip host 192.168.3.100 any

route-map GW_MAC_IP_TRACKING permit 10
 match ip address gatewayiplist
 match evpn route-type 2-mac-ip
 set extcommunity default-gw

    ! Attach the route-map to the neighbor

router bgp 65001
bgp router-id interface Loopback0 bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor
 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0
!
address-family ipv4 exit-address-family
!
address-family l2vpn evpn neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community both
neighbor 172.16.255.1 route-map GW_MAC_IP_TRACKING out
neighbor 172.16.255.2 activate
neighbor 172.16.255.2 send-community both
neighbor 172.16.255.2 route-map GW_MAC_IP_TRACKING out
exit-address-family
```



CHAPTER 8

Configuring VXLAN-Aware Flexible Netflow

- Restrictions for VXLAN-Aware Flexible NetFlow, on page 187
- Information About VXLAN-Aware Flexible NetFlow, on page 187
- How to Configure VXLAN-Aware Flexible NetFlow, on page 187
- Configuration Examples for VXLAN-Aware Flexible NetFlow, on page 192

Restrictions for VXLAN-Aware Flexible NetFlow

Information About VXLAN-Aware Flexible NetFlow

Flexible NetFlow (FNF) uses flows to provide statistics for accounting, network monitoring, and network planning. VXLAN-aware FNF provides information about the VXLAN-encapsulated IPV4 and IPV6 packets in the network. VXLAN-aware FNF captures the VXLAN flow information for both bridged and routed traffic.

A flow is a unidirectional stream of packets that arrives on a source interface and has the same values for the keys. A key is an identified value for a field within the packet. You create a flow using a flow record to define the unique keys for your flow. FNF allows you to define an optimal flow record for a particular application by selecting the keys from a large collection of predefined fields. All key values must match for the packet to count in a given flow. Flows are stored in the FNF cache. You can export the data FNF gathers for your flow by using an exporter.

In a BGP EVPN VXLAN fabric, an FNF monitor is configured on the NVE interface on a VTEP and on the physical interface on a spine switch. For more information about FNF, see *Configuring Flexible NetFlow* module of the *Network Management Configuration Guide*.

How to Configure VXLAN-Aware Flexible NetFlow

To configure VXLAN-aware FNF, perform these steps:

1. Create a flow record by specifying key fields and non-key fields to the flow.
2. Create a flow exporter by specifying the export protocol and transport destination port, source, and other parameters.
3. Create a flow monitor based on the flow record and flow exporter.

4. Apply the flow monitor to the network virtualization edge (NVE) interface on the VTEPs.



Note The commands listed in this section are applicable only to VXLAN-aware FNF. For detailed steps to configure FNF, see *How to Configure Flexible Netflow* section in the *Configuring Flexible NetFlow* module of the *Network Management Configuration Guide*.

Configuring a Flow Record

To configure a flow record for VXLAN-aware FNF, perform the following steps:



Note All the **match** commands listed in this configuration task are mandatory.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**flow record** *flow-record-name* **Example:** `Device(config)# flow record vxlan_nf_record_input`	Creates a flow record and enters flow record configuration mode. This command also allows you to modify an existing flow record. **Note** We recommend that you configure a unique flow record for each address family (IPv4 and IPv6) and also for each traffic direction (input and output). Ensure that the flow record for ingress traffic has the **match** commands configured with the **input** keyword. Ensure that the flow record for egress traffic has the **match** commands configured with the **output** keyword.

	Command or Action	Purpose
Step 4	**match datalink vlan** {**input** | **output**} **Example:** `Device(config-flow-record)# match datalink vlan output`	Configures the VLAN ID (for input or output traffic) as a key field for the FNF flow record. **Note** Ensure that you configure the **vlan input** and **vlan output** fields. These fields are required for VXLAN-aware FNF to work on EVPN input and output traffic flows.
Step 5	**match routing vrf input** **Example:** `Device(config-flow-record)# match routing vrf input`	Configures the VRF ID (for input or output traffic) as a key field for the FNF flow record. **Note** Ensure that you configure the **vrf input** field. This field is required for VXLAN-aware FNF to work on EVPN input and output traffic flows.
Step 6	**match vxlan vtep** {**input** | **output**} **Example:** `Device(config-flow-record)# match vxlan vtep output`	Configures the VTEP ID as a key field for the FNF flow record. The **input** keyword shows the VTEP source IP address in the captured flow. The **output** keyword shows the VTEP destination IP address in the captured flow
Step 7	**match vxlan vnid** **Example:** `Device(config-flow-record)# match vxlan vnid`	Configures the VXLAN VNI ID as a key field for the FNF flow record.
Step 8	**end** **Example:** `Device(config-flow-record)# end`	Returns to privileged EXEC mode.

Configuring Flow Exporter

To configure flow exporter for VXLAN-aware FNF, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**flow exporter** *flow-exporter-name* **Example:** `Device(config)# flow exporter e1`	Creates a flow exporter and enters flow exporter configuration mode.
Step 4	**destination** *ipv4-address* **Example:** `Device(config-flow-exporter)# destination 172.16.103.2`	Sets the IPv4 destination address or hostname for the exporter.
Step 5	**source** *interface-type interface-number* **Example:** `Device(config-flow-exporter)# source TenGigabitEthernet1/5/0/3`	Specifies the interface to use to reach the NetFlow collector at the configured destination. Ensure that the source IP address is unique per fabric. **Note** We recommend that you configure a unique Loopback on each VTEP. **Note** Flow exporter does not support unnumbered IP interface as source interface.
Step 6	**ttl** *seconds* **Example:** `Device(config-flow-exporter)# ttl 4`	Configures the time-to-live (TTL) value for datagrams sent by the exporter. The range is from 1 to 255 seconds. The default is 255.
Step 7	**transport udp** *port-number* **Example:** `Device(config-flow-exporter)# transport udp 2055`	Specifies the UDP port to use to reach the NetFlow collector.
Step 8	**export-protocol** {**ipfix** | **netflow-v9**} **Example:** `Device(config-flow-exporter)# export-protocol ipfix`	Specifies the version of the NetFlow export protocol used by the exporter.
Step 9	**end** **Example:** `Device(config-flow-exporter)# end`	Returns to privileged EXEC mode.

Configuring a Flow Monitor

To configure a flow monitor for VXLAN-aware FNF, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**flow monitor** *flow-monitor-name* **Example:** `Device(config)# flow monitor vxlan_nf_monitor_input`	Creates a flow monitor and enters flow monitor configuration mode. This command also allows you to modify an existing flow monitor.
Step 4	**exporter** *flow-exporter-name* **Example:** `Device(config-flow-monitor)# exporter e1`	Specifies the name of the flow exporter that was created previously and associates it with the specified flow monitor.
Step 5	**record** *flow-record-name* **Example:** `Device(config-flow-monitor)# record vxlan_nf_record_input`	Specifies the record for the flow monitor.
Step 6	**end** **Example:** `Device(config-flow-monitor)# end`	Returns to privileged EXEC mode.

Configuring Flexible NetFlow on an NVE Interface

To configure VXLAN-aware FNF on the NVE interface of a VTEP, perform the following steps:

Procedure

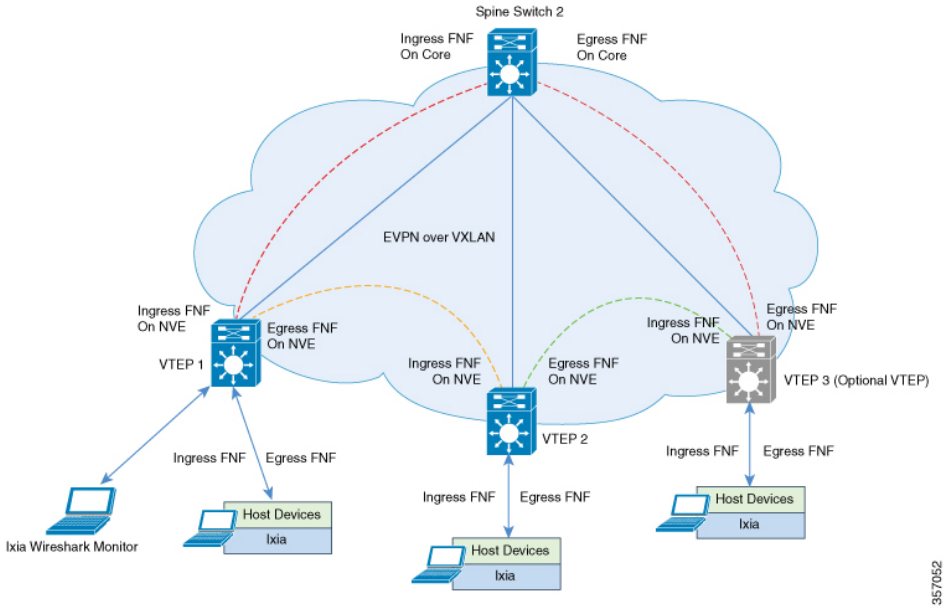
	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *nve-interface-number* **Example:** `Device(config)# interface nve1`	Specifies the network virtualization edge (NVE) interface number and enters interface configuration mode.
Step 4	**ip flow monitor** *flow-monitor-name* {**input** | **output**} **Example:** `Device(config-if)# ip flow monitor vxlan_nf_monitor_input input`	Associates the IPv4 flow monitor to the NVE interface for input or output packets.
Step 5	**ipv6 flow monitor** *flow-monitor-name* {**input** | **output**} **Example:** `Device(config-if)# ipv6 flow monitor vxlan_nf_v6monitor_input input`	Associates the IPv6 flow monitor to the NVE interface for input or output packets.
Step 6	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuration Examples for VXLAN-Aware Flexible NetFlow

This section provides configuration examples for VXLAN-aware FNF using the following topology:

Figure 19: EVPN VXLAN Topology with VXLAN-Aware Flexible NetFlow



Configuring VTEP 1 to enable VXLAN-Aware Flexible NetFlow

The following table provides a sample configuration for VTEP 1 to enable VXLAN-aware FNF:

Table 26: Configuring VTEP 1 to enable VXLAN-Aware Flexible NetFlow

VTEP 1

```
Leaf-01# show running-config

<snip: only config relevant to vxlan netflow is shown>
flow record vxlan_nf_record_input
 match datalink vlan input
 match datalink mac source address input
 match datalink mac destination address input
 match routing vrf input
 match ipv4 ttl
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 match transport icmp ipv4 type
 match transport icmp ipv4 code
 match transport igmp type
 match interface input
 match flow direction
 match vxlan vnid
 match vxlan vtep input
 match vxlan vtep output
 collect counter bytes long
 collect counter packets long
 collect timestamp absolute first
 collect timestamp absolute last
!
flow record vxlan_nf_record_output
 match datalink mac destination address output
 match ipv4 protocol
 match ipv4 source address
 match ipv4 destination address
 match transport source-port
 match transport destination-port
 match datalink vlan output
 match vxlan vnid
 match vxlan vtep input
 match vxlan vtep output
 collect counter bytes long
 collect counter packets long
 collect timestamp absolute first
 collect timestamp absolute last
!
```

VTEP 1

```
flow record vxlan_nf_v6record_input
 match datalink vlan input
 match routing vrf input
 match ipv6 protocol
 match ipv6 source address
 match ipv6 destination address
 match transport source-port
 match transport destination-port
 match vxlan vnid
 match vxlan vtep input
 match vxlan vtep output
 collect counter bytes long
 collect counter packets long
 collect timestamp absolute first
 collect timestamp absolute last
!
flow record vxlan_nf_v6record_output
  match datalink vlan output
 match ipv6 protocol
 match ipv6 source address
 match ipv6 destination address
 match transport source-port
 match transport destination-port
 match vxlan vnid
 match vxlan vtep input
 match vxlan vtep output
 collect counter bytes long
 collect counter packets long
 collect timestamp absolute first
 collect timestamp absolute last
!
flow exporter e1
 destination 172.16.103.2
 source TenGigabitEthernet1/5/0/3
 ttl 4
 transport udp 2055
 export-protocol ipfix
!
flow monitor vxlan_nf_monitor_input
 exporter e1
 cache timeout inactive 100
 cache timeout active 100
 record vxlan_nf_record_input
 !
!
flow monitor vxlan_nf_monitor_output
 exporter e1
 cache timeout inactive 100
 cache timeout active 100
 record vxlan_nf_record_output
 !
!
flow monitor vxlan_nf_v6monitor_input
 exporter e1
 cache timeout inactive 100
 cache timeout active 100
 record vxlan_nf_v6record_input
 !
!
```

VTEP 1

```
flow monitor vxlan_nf_v6monitor_output
 exporter e1
 cache timeout inactive 100
 cache timeout active 100
 record vxlan_nf_v6record_output
!
interface nve1
 ip flow monitor vxlan_nf_monitor_input input
 ip flow monitor vxlan_nf_monitor_output output
 ipv6 flow monitor vxlan_nf_v6monitor_input input
 ipv6 flow monitor vxlan_nf_v6monitor_output output
!
Leaf-01#
```

Checking IPv4 Input Flow Monitor Cache Output

The following example provides a sample output to check the IPv4 input flow monitor cache output on VTEP 1:

```
Leaf-01# configure terminal
Leaf-01(config)# show flow monitor vxlan_nf_monitor_input cache format table

  Cache type:                               Normal (Platform cache)

  Cache size:                                10000

  Current entries:                               4



  Flows added:                                   8

  Flows aged:                                    4

    - Inactive timeout    (   100 secs)          4
DATALINK VLAN INPUT  DATALINK MAC SRC ADDR INPUT  DATALINK MAC DST ADDR INPUT  IP VRF ID
INPUT                IPV4 SRC ADDR    IPV4 DST ADDR    TRNS SRC PORT  TRNS DST PORT  ICMP
IPV4 TYPE  ICMP IPV4 CODE  IGMP TYPE  INTF INPUT           FLOW DIRN      VXLAN VXLAN VNID
  VXLAN VXLAN VTEP INPUT  VXLAN VXLAN VTEP OUTPUT  IP PROT  IP TTL            bytes long
          pkts long  time abs first  time abs last
==================  ==========================  ==========================
============================  ===============  ===============  =============  =============
  ==============  ==============  =========  ====================  =========
====================  ======================  ======================  =======  ======
====================  ====================  ==============  ==============
                 13  AAAA.CCCC.1003               AAAA.BBBB.1003               3
(l3vni5001)         192.168.13.3     192.168.13.2                 0              0
       0               0          0  Null                  Input                    10013
  2.2.2.2                 1.1.1.1                        61      64              43517376
               43172    14:00:41.391   14:01:34.391
                 11  AAAA.CCCC.1001               AAAA.BBBB.1001               2
(l3vni5000)         192.168.11.3     192.168.11.2                 0              0
       0               0          0  Null                  Input                    10011
  2.2.2.2                 1.1.1.1                        61      64              43517376
               43172    14:00:41.391   14:01:34.391
                 10  AAAA.CCCC.1002               AAAA.BBBB.1002               2
(l3vni5000)         192.168.10.3     192.168.10.2                 0              0
       0               0          0  Null                  Input                    10010
  2.2.2.2                 1.1.1.1                        61      64              43517376
               43172    14:00:41.391   14:01:34.391
```

```
                12  AAAA.CCCC.1004               AAAA.BBBB.1004               3
(l3vni5001)         192.168.12.3     192.168.12.2                0              0
       0               0          0  Null                  Input                    10012
  2.2.2.2                 1.1.1.1                       61      64              43517376
               43172    14:00:41.391   14:01:34.391

Leaf-01#
```

Checking IPv4 Output Flow Monitor Cache Output

The following example provides a sample output to check the IPv4 output flow monitor cache output on VTEP 1:

```
Leaf-01# configure terminal
Leaf-01(config)# show flow monitor vxlan_nf_monitor_output cache format table

  Cache type:                               Normal (Platform cache)

  Cache size:                                10000

  Current entries:                               4

  Flows added:                                   8

  Flows aged:                                    4

    - Inactive timeout    (   100 secs)          4

DATALINK MAC DST ADDR OUTPUT  IPV4 SRC ADDR    IPV4 DST ADDR    TRNS SRC PORT  TRNS DST
PORT  DATALINK VLAN OUTPUT      VXLAN VXLAN VNID  VXLAN VXLAN VTEP INPUT  VXLAN VXLAN VTEP
 OUTPUT  IP PROT           bytes long             pkts long  time abs first  time abs last
============================  ===============  ===============  =============  =============
  ====================  ====================  ======================  ======================
  =======  ====================  ====================  ==============  =============
AAAA.CCCC.1002                192.168.10.2     192.168.10.3                0
 0                   10                10010  1.1.1.1                2.2.2.2
            61              44812536                43172    14:00:41.391   14:01:34.391
AAAA.CCCC.1004                192.168.12.2     192.168.12.3                0
 0                   12                10012  1.1.1.1                2.2.2.2
            61              44812536                43172    14:00:41.391   14:01:34.391
AAAA.CCCC.1003                192.168.13.2     192.168.13.3                0
 0                   13                10013  1.1.1.1                2.2.2.2
            61              44812536                43172    14:00:41.391   14:01:34.391
AAAA.CCCC.1001                192.168.11.2     192.168.11.3                0
 0                   11                10011  1.1.1.1                2.2.2.2
            61              44812536                43172    14:00:41.391   14:01:34.391

Leaf-01#
```

Checking IPv6 Input Flow Monitor Cache Output

The following example provides a sample output to check the IPv6 input flow monitor cache output on VTEP 1:

```
Leaf-01# configure terminal
Leaf-01(config)# show flow monitor vxlan_nf_v6monitor_input cache format table

  Cache type:                               Normal (Platform cache)

  Cache size:                                10000
```

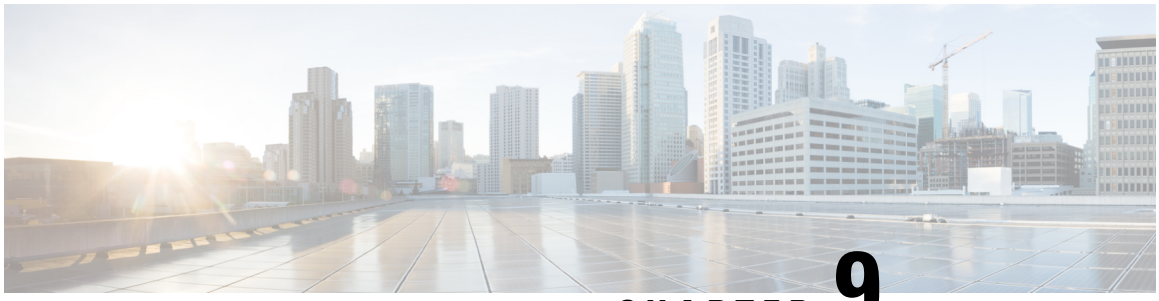
```
  Current entries:                              4


  Flows added:                                  8

  Flows aged:                                   4

    - Inactive timeout    (   100 secs)         4
IPV6 SRC ADDR                                IPV6 DST ADDR
    TRNS SRC PORT  TRNS DST PORT      VXLAN VXLAN VNID  VXLAN VXLAN VTEP INPUT  VXLAN VXLAN
 VTEP OUTPUT  IP PROT            bytes long             pkts long  time abs first  time abs
 last
============================================  ============================================
  =============  =============  ====================  ======================
======================  =======  ====================  ====================  ==============
  =============
192:168:12::3                                192:168:12::2
                 0              0                10012  2.2.2.2                1.1.1.1
                     59              43517376                43172    14:00:41.391
14:01:34.391
192:168:10::3                                192:168:10::2
                 0              0                10010  2.2.2.2                1.1.1.1
                     59              43517376                43172    14:00:41.391
14:01:34.391
192:168:13::3                                192:168:13::2
                 0              0                10013  2.2.2.2                1.1.1.1
                     59              43517376                43172    14:00:41.391
14:01:34.391
192:168:11::3                                192:168:11::2
                 0              0                10011  2.2.2.2                1.1.1.1
                     59              43517376                43172    14:00:41.391
14:01:34.391

Leaf-01#
```



CHAPTER 9

Configuring BGP EVPN VXLAN QoS

BGP EVPN VXLAN QoS enables you to provides Quality of Service (QoS) capabilities to traffic that is EVPN VXLAN-encapsulated.

By configuring QoS you can provide preferential treatment to specific types of traffic at the expense of other traffic types. Without QoS, the device offers best-effort service for each packet, regardless of the packet contents or size; the device sends the packets without any assurance of reliability, delay bounds, or throughput.

QoS provides the following features:

- Low latency
- Bandwidth guarantee
- Buffering capabilities and dropping disciplines
- Traffic policing
- Enables the changing of the attribute of the frame or packet header
- Relative services

For more information about QoS and its configuration, see *Quality of Service Configuration Guide, Cisco IOS XE* for the device.

- Quality of Service Components, on page 199
- How to Configure BGP EVPN VXLAN QoS, on page 200
- BGP EVPN VXLAN QoS Modes, on page 201
- BGP EVPN VXLAN Traffic without QoS Policy , on page 202
- BGP EVPN VXLAN QoS Marking with an Input Policy, on page 203
- BGP EVPN VXLAN QoS Marking with Input and Output Policies, on page 204

Quality of Service Components

QoS consists of the following key components:

QoS Component	Description
Classification	Classification is the process of distinguishing one type of traffic from another based upon access control lists (ACLs), Differentiated Services Code Point (DSCP), Class of Service (CoS), and other factors.

QoS Component	Description
Marking	Marking is used on traffic to convey specific information to a downstream device in the network, or to carry information from one interface in a device to another. When traffic is marked, QoS operations on that traffic can be applied. This can be accomplished directly using the set command or through a table map, which takes input values and translates them directly to values on output. Through marking, you can set: a Layer 2 COS value in a frame, Layer 3 DSCP value in a packet, and Layer 3 ECN value in a packet.
Shaping and Policing	Shaping is the process of imposing a maximum rate of traffic, while regulating the traffic rate in such a way that downstream devices are not subjected to congestion. Shaping in the most common form is used to limit the traffic sent from a physical or logical interface. Policing is used to impose a maximum rate on a traffic class. If the rate is exceeded, then a specific action is taken as soon as the event occurs.
Queuing	Queuing is used to prevent traffic congestion. Traffic is sent to specific queues for servicing and scheduling based upon bandwidth allocation. Traffic is then scheduled or sent out through the port.
Bandwidth	Bandwidth allocation determines the available capacity for traffic that is subject to QoS policies.
Trust	Trust enables traffic to pass through the device, and the Differentiated Services Code Point (DSCP), precedence, or CoS values coming in from the end points are retained in the absence of any explicit policy configuration.
Class of Service (CoS)	CoS refers to the three bits in an 802.1Q header that are used to indicate the priority of the Ethernet frame as it passes through a switched network. The CoS bits in the 802.1Q header are commonly referred to as the 802.1p bits. 802.1Q is discarded prior to frame encapsulation in a VXLAN header, where CoS value is not present in VXLAN tunnel. To maintain QoS when a packet enters the VXLAN tunnel, the type of service (ToS) and CoS values map to each other.
Differentiated Services Code Point (DSCP)	The first six bits of the ToS byte in the IP header. DSCP is only present in an IP packet.
Explicit Congestion Notification (ECN)	The last two bits of the ToS byte in the IP header. ECN is only present in an IP packet.

How to Configure BGP EVPN VXLAN QoS

QoS features in a BGP EVPN VXLAN fabric are enabled through the Modular QoS CLI (MQC). The MQC is a CLI structure that allows you to create traffic policies and attach these policies to interfaces. A traffic policy contains a traffic class and one or more QoS features. A traffic class is used to classify traffic, while the QoS features in the traffic policy determine how to treat the classified traffic.

MQC uses the following three steps to configure QoS:

1. Define a traffic class with the **class-map** command.

2. Create a traffic policy by associating the traffic class with one or more QoS features, using the **policy-map** command.
3. Attach the traffic policy to the interface with the **service-policy** command.

For more information about configuring QoS, see the *Quality of Service Configuration Guide, Cisco IOS XE* for the device.

Note In a BGP EVPN VXLAN fabric, QoS policies are applied on the physical interface of the VTEPs in the underlay. The NVE interface of the VTEPs support only the BUM rate limiter policies. For a sample configuration of BUM Traffic Rate Limiting, see "Configuring EVPN VXLAN Layer 2 Overlay Network" chapter in the *BGP EVPN VXLAN Configuration Guide* for the relevant release.

BGP EVPN VXLAN QoS Modes

Layer 2 or Layer 3 traffic enters a BGP EVPN VXLAN fabric at its edge node, which is also known as the ingress VTEP. The ingress VTEP encapsulates the original packet in a VXLAN header. The VXLAN header of the encapsulated packet is assigned a DSCP value based on QoS rules. The Spines and any other intermediate Layer 3 network devices are transit systems interconnecting the ingress and egress VTEPs. These intermediate devices use the outer header of the encapsulated packet to perform classification, marking, and policing. The encapsulated packet exits the fabric through the egress VTEP. At the egress VTEP, the VXLAN encapsulated packets are decapsulated and Ethernet packets are sent to the destination. The DSCP value of the inner header (decapsulated packet header) is used to perform classification, marking, and policing.

EVPN VXLAN QoS is supported in two modes:

Ingress VTEP: Uniform Mode

When a packet enters the fabric, if the ingress interface has a QoS policy that first marks the original packet and then the packet is encapsulated, both the inner and outer headers of the packet have the same updated DSCP value.

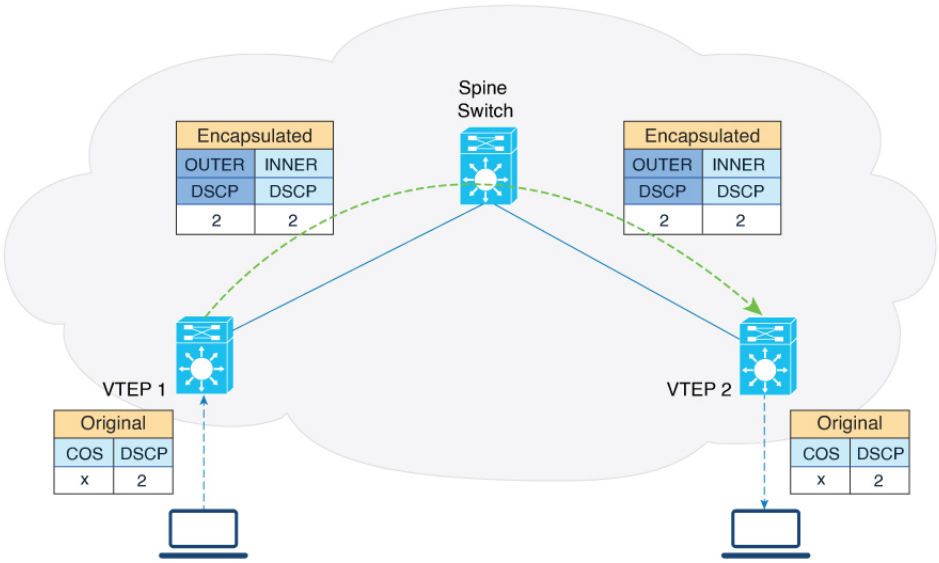
Uniform mode is the default mode of QoS operation.

Egress VTEP: Pipe Mode

As the encapsulated packet traverses through the fabric, if an intermediate node has a QoS policy that alters the outer DSCP value, the outer header is removed through decapsulation at egress. The DSCP value from the outer header is not propagated to the original packet. In this way, the original ToS value is preserved at the end of the VXLAN tunnel.

BGP EVPN VXLAN Traffic without QoS Policy

Figure 20: EVPN VXLAN Traffic without a QoS Policy



Encapsulation at Ingress VTEP

1. Traffic entering the edge of EVPN VXLAN fabric is received by the ingress VTEP (Leaf switch VTEP1) on its access interface.
2. Ingress VTEP encapsulates the traffic packets with a VXLAN header. During the VXLAN encapsulation, the DSCP value from inner header is copied to the outer header.

Transportation of a packet through the VXLAN Tunnel

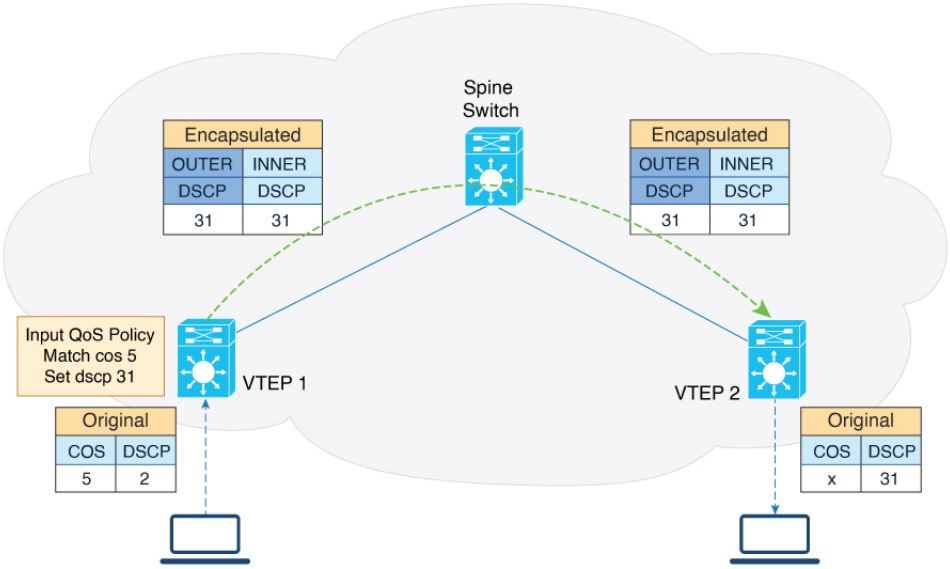
As the VXLAN encapsulated packet traverses through the EVPN VXLAN fabric, the intermediate switches use the outer header to perform classification, marking, and policing.

Decapsulation at the Egress VTEP

1. When a VXLAN-encapsulated packet reaches the NVE interface of the egress VTEP (Leaf switch VTEP2), the switch removes the VXLAN header from the packet.
2. It uses the DSCP value of the inner header to perform classification, marking, and policing.
3. The DSCP value from the outer header is not propagated into the original packet header if the outer DSCP value was changed by any intermediate switch.
4. If a packet exits on a trunk port, the COS value for the packet is derived from the DSCP-to-COS table map.

BGP EVPN VXLAN QoS Marking with an Input Policy

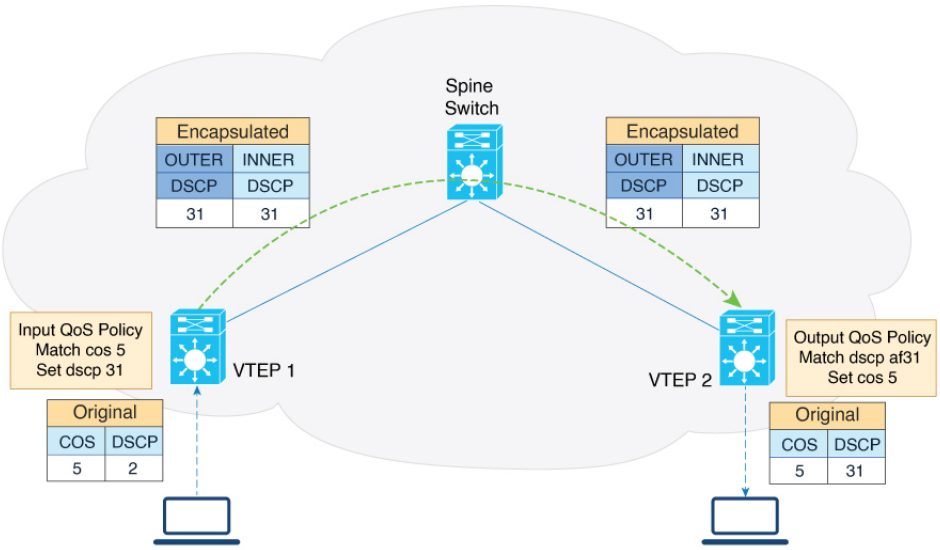
Figure 21: EVPN QoS Marking with an Input Policy



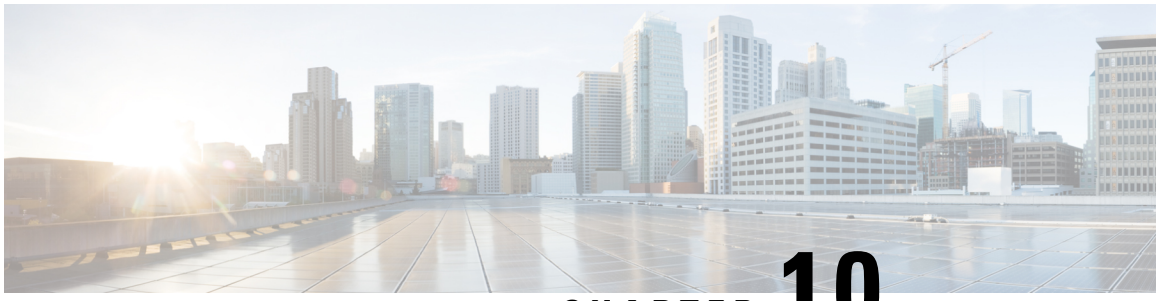
1. The ingress VTEP VXLAN-encapsulates the packet and marks both the outer and inner header. During the VXLAN encapsulation, the IP packet's DSCP value is modified based on the input policy.
2. At egress, the outer VXLAN header is removed and the outer header DSCP values are not propagated to the decapsulated packets.

BGP EVPN VXLAN QoS Marking with Input and Output Policies

Figure 22: EVPN QoS Marking with an Input and an Output Policy



1. At ingress, the VTEP encapsulates the packet and marks both the outer and inner header. During the VXLAN encapsulation, the original packet's DSCP value is modified based on the input policy.
2. At egress, the outer VXLAN header is removed and the outer header DSCP or COS values are marked based on the egress QoS policy.



CHAPTER 10

Configuring Optimized Layer 2 Overlay Multicast

- Restrictions for Optimized Layer 2 Overlay Multicast, on page 205
- Information About Optimized Layer 2 Overlay Multicast, on page 205
- How to Configure Optimized Layer 2 Overlay Multicast, on page 209
- Verifying Optimized Layer 2 Overlay Multicast on VTEP and Centralized Gateway, on page 213
- Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214

Restrictions for Optimized Layer 2 Overlay Multicast

- This feature does not support data mutlicast distribution tree (MDT) in the underlay.
- Optimized Layer 2 overlay source specific multicast forwarding (S,G) is not supported. Only (*,G) is supported.
- Optimized Layer 2 overlay multicast is ineffective for a Distributed Anycast Gateway deployment (use Layer 3 tenant routed multicast in a Distributed Anycast Gateway deployment). Optimized Layer 2 overlay multicast is applicable between Layer 2 leaf or the centralized gateway that extends the bridge-domain over the EVPN VXLAN fabric.
- Optimized Layer 2 overlay multicast handoff to Layer 3 tenant routed multicast (TRM) within the EVPN fabric is not supported.

Information About Optimized Layer 2 Overlay Multicast

Optimized Layer 2 Overlay Multicast functionality is based on Internet Group Management Protocol (IGMP) or Multicast Listener Discovery (MLD) Proxy for EVPN, as defined in draft-ietf-bess-evpn-igmp-mld-proxy.

In a BGP EVPN VXLAN fabric, there can be many hosts attached to a VLAN that is stretched across several sites. These hosts send IGMP membership reports to join their interested multicast groups. Also, the IGMP router periodically sends membership queries to find out if there are hosts on a VLAN that are interested in receiving multicast traffic for that group. This leads to a flooding of IGMP reports and queries in the overlay network across multiple sites. The network may have few hosts, but the IGMP messages are sent to all VTEPs, resulting in an overloading of the fabric. This necessitates an IGMP or MLD Proxy to efficiently forward the multicast traffic. When you enable an IGMP or MLD Proxy on the VTEP, the proxy forwards the multicast traffic only to the interested receivers, thus optimizing the bandwidth of the fabric.

Optimized Layer 2 Overlay Multicast functionality facilitates multicast traffic both at the access level and within the VXLAN fabric. Optimized Layer 2 Overlay Multicast is restricted to a Layer 2 Virtual Network Instance (L2VNI), wherein the source and receivers are within the same Layer 2 domain. Layer 2 multicast data traffic is not forwarded to a VTEP or a port that does not have a multicast router or receiver attached to it.

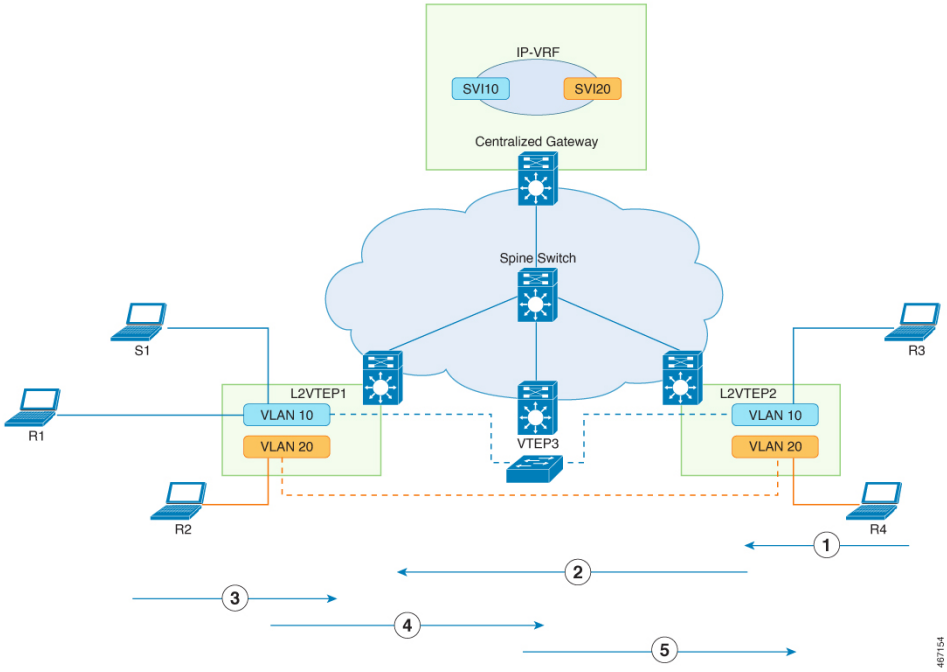
Optimized Layer 2 Overlay Multicast supports both IPv4 and IPv6 multicast traffic in a Layer 2 EVPN VXLAN fabric.

Optimized Layer 2 Multicast in BGP EVPN VXLAN Layer 2 Overlay

Optimized Layer 2 Overlay Multicast uses an EVPN instance to build an IGMP or MLD overlay. When a VTEP receives an IGMP or MLD report from a host on the access link, it propagates this information to all other VTEPs in the same EVPN instance, using EVPN Route Type 6 (RT6). Propagation of RT6 updates the multicast forwarding tree for the intra-VLAN multicast traffic in the fabric. If a host on a second VTEP sends multicast traffic to a group that already has receivers, traffic is forwarded in the established multicast forwarding tree. Multicast traffic source and receiver, which are on the same VLAN, but stretched across VTEPs in the BGP EVPN VXLAN fabric, can communicate with each other.

Source and Receiver on Same VLAN, Connected to Different Layer 2 VTEPs

Figure 23: Topology Showing Layer 2 Overlay Multicast Traffic Within a VLAN



1. Receiver R3 in VLAN 10 sends (*,G) IGMP Join to VTEP 2. (*,G) is created at VTEP 2.
2. VTEP 2 sends EVPN route type 6 to VTEP 1. (*,G) is created at VTEP 1.
3. Source S1 in VLAN 10 starts streaming data. VTEP 1 forwards the multicast packets to all other VTEPs that have receivers.

4. Layer 2 multicast packets are encapsulated with L2VNI that corresponds to VLAN 10 and are forwarded through the overlay to VTEP 2.
5. On VTEP 2, the L2VNI encapsulated data packets are decapsulated and forwarded to the port where receiver R3 is connected.

Optimized Layer 2 Overlay Multicast Interworking in BGP EVPN VXLAN Fabric

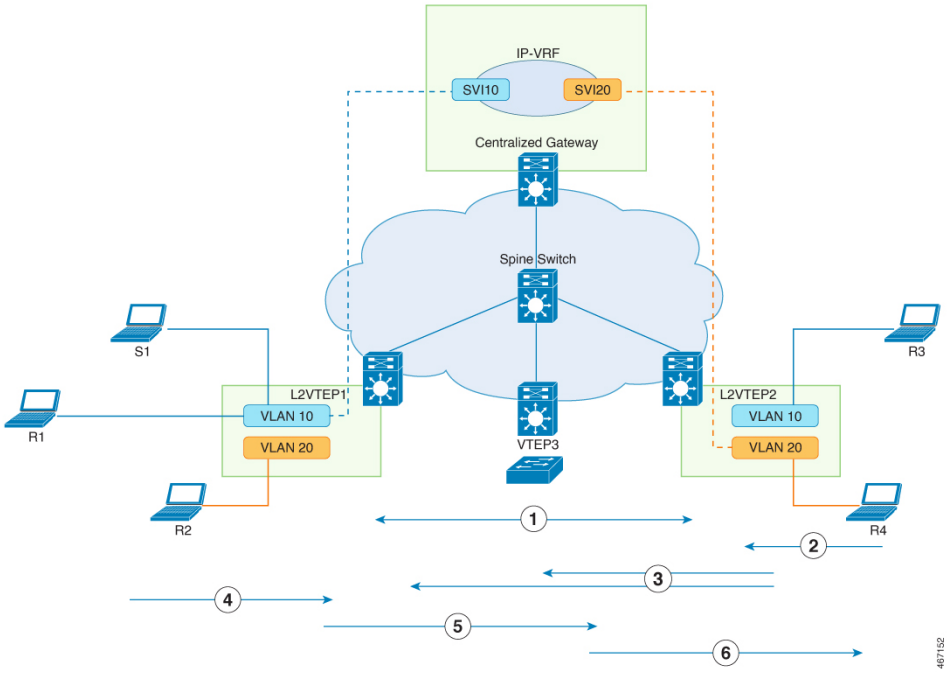
VTEPs in the EVPN VXLAN network forward traffic to each other through the VXLAN gateways. An EVPN VXLAN Centralized Gateway (CGW) VTEP performs the Layer 3 gateway function for all the Layer 2 virtual network instances (VNIs). All the other VTEPs in the network perform only bridging. For information on CGW, see Configuring EVPN VXLAN Integrated Routing and Bridging.

Optimized Layer 2 Multicast Across Layer 2 Overlays

In an EVPN VXLAN fabric, a CGW performs routing for the inter-L2VNI Layer 2 multicast traffic. In this case, the sender and receiver, can both be on the same VTEP or on different VTEPs within the EVPN VXLAN fabric.

Source and Receiver on Different VLANs, Connected to Different Layer 2 VTEPs

Figure 24: Topology Showing Layer 2 Overlay Multicast Traffic Between Different VLANs



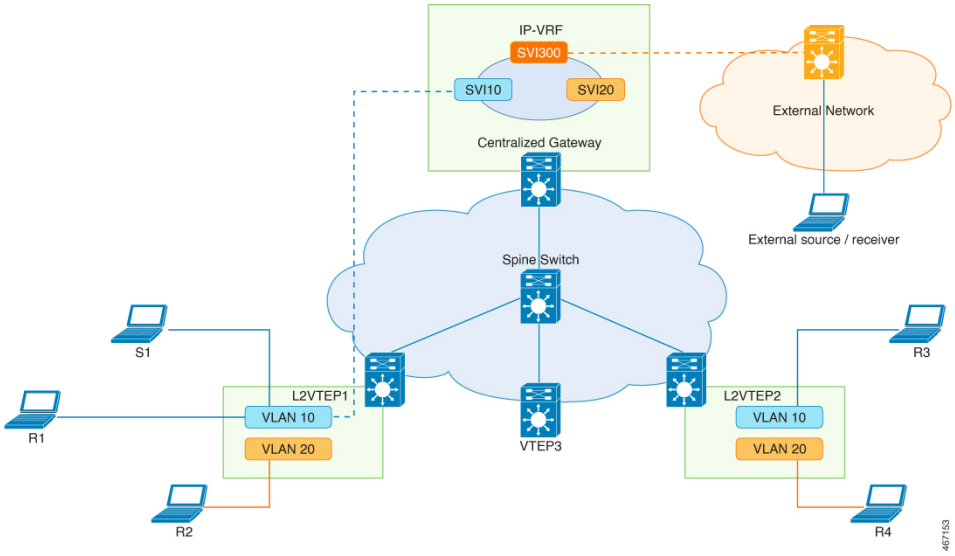
1. CGW sends EVPN route type 6 (*, *) to all the Layer 2 VTEPs in EVPN VXLAN fabric.
2. Receiver R4 in VLAN 20 sends (*, G) IGMP join to VTEP 2. (*, G) is created in IGMP snooping table on VTEP 2.
3. VTEP 2 sends EVPN route type 6 (*, G) to the CGW and all other Layer 2 VTEPs. (* , G) is also created in Layer 3 multicast on the CGW.

4. Source S1 in VLAN 10 starts streaming data. VTEP1 forwards the multicast traffic.
5. Layer 2 multicast traffic is switched through L2VNI to the CGW, based on the default EVPN route type 6.

   In the CGW, multicast data packets are processed by Layer 3 Multicast and routed to access SVI20, where the receiver R4 is connected.
6. Multicast traffic is forwarded to VTEP 2 through L2VNI. On egress VTEP 2, Layer 2 Overlay Multicast switches the data packets to the port where the receiver R4 is connected.

Optimized Layer 2 Overlay Multicast Handoff with External Network

Figure 25: Topology Showing Layer 2 Overlay Multicast Traffic Handoff to an External Network



In this topology, either the multicast receiver or sender exists on an external network. The Centralized Gateway is configured to route the multicast traffic between the EVPN VXLAN fabric and the external network. The leaf VTEPs in the VXLAN fabric perform only bridging.

If a source device that is connected to a Layer 2 VTEP starts streaming multicast traffic, the traffic is bridged from the Layer 2 VTEP to the Centralized Gateway VTEP. The Centralized Gateway VTEP then routes the multicast traffic to the interested receivers in the external network.

The rendezvous point (RP) is located outside the fabric.

Replication Types Supported in an Underlay Network

Optimized Layer 2 Overlay Multicast (OL2OM) supports the following replication types:

- Ingress Replication
- Underlay Multicast with Default MDT

Ingress Replication

Ingress replication, or headend replication, is a unicast approach to handle multidestination Layer 2 overlay broadcast, unknown unicast, and multicast (BUM) traffic. OL2OM uses the ingress replication mechanism to send the Layer 2 overlay multicast traffic only to those VTEPs where receivers have joined the multicast group. If ingress replication is already configured as part of the initial Layer 2 VNI configuration, OL2OM uses the same mode of replication, without the need for extra configuration steps. Using ingress replication to handle OL2OM traffic can result in scaling issues because an ingress device must replicate the multicast traffic as many times as the number of VTEPs that are associated with the Layer 2 VNI. Use ingress replication in a network that has fewer receivers.

Underlay Multicast with Default MDT

A multicast distribution tree (MDT) is a multicast tunnel that transports multicast traffic. Traffic through the default MDT is flooded to all remote VTEPs, irrespective of whether the VTEP has any receivers or not. For more information, see Underlay Multicast . For a large-scale fabric network, we recommend Underlay Multicast Replication. Optimized Layer 2 Overlay Multicast traffic can take advantage of the underlying multicast network for data replication to support better scale and performance.

How to Configure Optimized Layer 2 Overlay Multicast

Configure Optimized Layer 2 Overlay Multicast in the fabric and also on the host-facing or Layer 2 access devices.

Prerequisites for Configuring Optimized Layer 2 Overlay Multicast

Before configuring Optimized Layer 2 Overlay Multicast, ensure that the EVPN VXLAN Layer 2 Overlay network has been configured by performing all the procedures that are listed in How to Configure EVPN VXLAN Layer 2 Overlay Network, on page 16.

Configure Optimized Layer 2 Multicast in EVPN VXLAN Layer 2 Overlay

The following sections provide information about how to configure optimized Layer 2 multicast in an EVPN VXLAN Layer 2 overlay.

Enable Optimized Layer 2 Overlay Multicast in BGP EVPN VXLAN Fabric

Follow these steps to enable Optimized Layer 2 Overlay Multicast in the EVPN VXLAN fabric.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:**	Enters global configuration mode.

	Command or Action	Purpose
	`Device# configure terminal`	
Step 3	**l2vpn evpn** **Example:** `Device(config)# l2vpn evpn`	Enters EVPN configuration mode.
Step 4	**multicast advertise** **Example:** `Device(config-evpn)# multicast advertise`	Enables Optimized Layer 2 multicast for all EVPN instances on the device. Optionally, you can enable or disable Optimized Layer 2 Multicast for an EVPN instance using the **multicast advertise** { **enable** | **disable** } command in the evpn-evi configuration mode. Example: `Device(config-evpn-evi)# multicast advertise disable`
Step 5	**end** **Example:** `Device(config-evpn)# end`	Returns to privileged EXEC mode.

Configure Optimized Layer 2 Overlay Multicast for IPv4

In addition to the steps provided in Prerequisites for Configuring Optimized Layer 2 Overlay Multicast, on page 209, perform the following steps to configure Optimized Layer 2 Overlay Multicast on the access devices that have IPv4 endpoints.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip igmp snooping querier** **Example:** `Device(config)# ip igmp snooping querier`	Enables the IGMP snooping querier.
Step 4	**ip igmp snooping querier version** *version-num* **Example:** `Device(config)# ip igmp snooping querier version 3`	Configures the IGMP version that the querier feature uses. By default, IGMP Version 2 is configured.

	Command or Action	Purpose
		IGMP Version 2 extends the IGMP functionality with features such as the IGMP leave process, group-specific queries, and an explicit maximum query response time. IGMP Version 3 provides the functionality for Basic IGMPv3 Snooping Support (BISS), which includes support for the snooping features and for IGMPv3 membership report messages.

Step 5	**ip igmp snooping querier** [**max-response-time** *seconds*] [**timer expiry** *seconds*]
	Example:
	`Device(config)# ip igmp snooping querier max-response-time 25`
	`Device(config)# ip igmp snooping querier timer expiry 200`
		(Optional) Configures the IGMP snooping querier parameters.
		max-response-time configures a maximum response time for snooping query messages. The range is from 1 to 25 seconds.
		timer expiry sets the length of time until the IGMP querier expires. The range is from 60 to 300 seconds.

Step 6	**end** **Example:** `Device(config-evpn)# end`	Returns to privileged EXEC mode.

Configure Optimized Layer 2 Overlay Multicast for IPv6

In addition to the information provided in the Prerequisites for Configuring Optimized Layer 2 Overlay Multicast, on page 209, perform the following steps to configure Optimized Layer 2 Overlay Multicast on the access devices that support IPv6 endpoints.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ipv6 mld snooping** **Example:** `Device(config)# ipv6 mld snooping`	Enables the MLD snooping.
Step 4	**ipv6 mld snooping querier** **Example:**	Enables the MLD snooping querier.

	Command or Action	Purpose
	`Device(config)# ipv6 mld snooping querier`	
Step 5	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configure Optimized Layer 2 Overlay Multicast Interworking

Optimized Layer 2 Overlay Multicast Interworking supports the following scenarios:

- Optimized Layer 2 Multicast across Layer 2 Overlays within the same fabric domain
- Optimized Layer 2 Overlay Multicast handoff with IP Multicast external network

Before You Begin

Perform the steps in the sequence listed here:

1. Before you configure the interworking scenarios, ensure that you set up the centralized gateway by performing all the steps described in Configuring EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway, on page 122 .
2. Enable Optimized Layer 2 Multicast on the centralized gateway by performing all the configurations described earlier in this document, under Configure Optimized Layer 2 Multicast in EVPN VXLAN Layer 2 Overlay.

Enable Layer 3 Multicast on Centralized Gateway

Follow these steps to enable Layer 3 Multicast on the Centralized Gateway.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	{**ip** | **ipv6**} **multicast-routing** **Example:** `Device(config)# ip multicast-routing`	Enables IP multicast routing. **Note** If the Centralized Gateway is associated with a VRF, run the **ip multicast-routing vrf** *vrf-name* command.

	Command or Action	Purpose
Step 4	**interface vlan** *vlan-id* **Example:** `Device(config)# interface vlan 20`	Specifies the interface on which you want to enable IP multicast and enters the interface configuration mode.
Step 5	**ip pim sparse-mode** **Example:** `Device(config-if)# ip pim sparse-mode`	Enables IP multicast on the interface.
Step 6	**ip igmp version** *version* **Example:** `Device(config-if)# ip igmp version 3`	Enables IGMP version on this interface. The default version of IGMP is set to Version 2. Repeat Step 4 to Step 6 for all the L2VNIs that have a source or receiver.
Step 7	**exit** **Example:** `Device(config-if)# exit`	Returns to global configuration mode.
Step 8	{**ip** | **ipv6**} **pim** [**vrf** *vrf-name*] **rp-address** *rp-address* **Example:** `Device(config)# ip pim rp-address 99.99.99.99`	Configures the Layer 3 Multicast rendezvous point (RP).

Verifying Optimized Layer 2 Overlay Multicast on VTEP and Centralized Gateway

Use the following **show** commands to verify the configuration of Optimized Layer 2 Overlay Multicast on VTEP and Centralized Gateway.

Command	Purpose
show ip igmp snooping	Verify that EVPN IGMP proxy is enabled in IGMP snooping.
show ip igmp snooping querier	Verify that the IGMP snooping querier version is as expected.
show ip igmp snooping groups	Verify the local and remote multicast group membership in IGMP snooping.
show ip mroute vrf <*vrf*>	Verify the MVPN routing and forwarding instance for the specified VRF.
show ip igmp vrf <*vrf-name*> groups detail	Verify the multicast groups with receivers that were learned through IGMP, on the centralized gateway.

Command	Purpose
show ipv6 mld snooping querier	Verify that the MLD snooping querier version is as expected.
show ipv6 mld snooping membership	Verify the local and remote multicast group membership in MLD Snooping.
show ipv6 mroute vrf <*vrf-name*> <*multicast-group-address*>	Verify the forwarding status of each IPv6 multicast route.
show ipv6 mld vrf <*vrf-name*> groups detail	Verify the multicast groups with receivers that were learned through MLD, on the centralized gateway.
show l2vpn evpn default-gateway	Verify the entries of the default gateway database.
show l2vpn evpn multicast [local] [remote]	Verify the local and remote multicast group membership in EVPN Manager.
show l2vpn evpn evi <*evpn-instance*> detail	Verify the Layer 2 multicast status for the specified EVPN instance.
show l2route evpn multicast smet [local] [remote] [local combined]	Verify the Selective Multicast Ethernet Tag Route (SMET) routes that the Layer 2 Route Information Base (L2RIB) receives from EVPN.
show l2route evpn multicast routes group <*multicast-address*>	Verify the multicast routes that are sent to L2FIB.
show bgp l2vpn evpn route-type 6	Verify the local and remote SMET routes in BGP.
show bgp l2vpn evpn summary	Verify the BGP EVPN routes on the spine switch.
show ip pim neighbor	Verify the Protocol Independent Multicast (PIM) neighbors that are discovered by the device.
show ip pim rp mapping	Verify that the PIM group-to-rendezvous point (RP) mappings are populated correctly on the device.
show ip msdp peer	Verify the Multicast Source Discovery Protocol (MSDP) peer information on the spine switch.

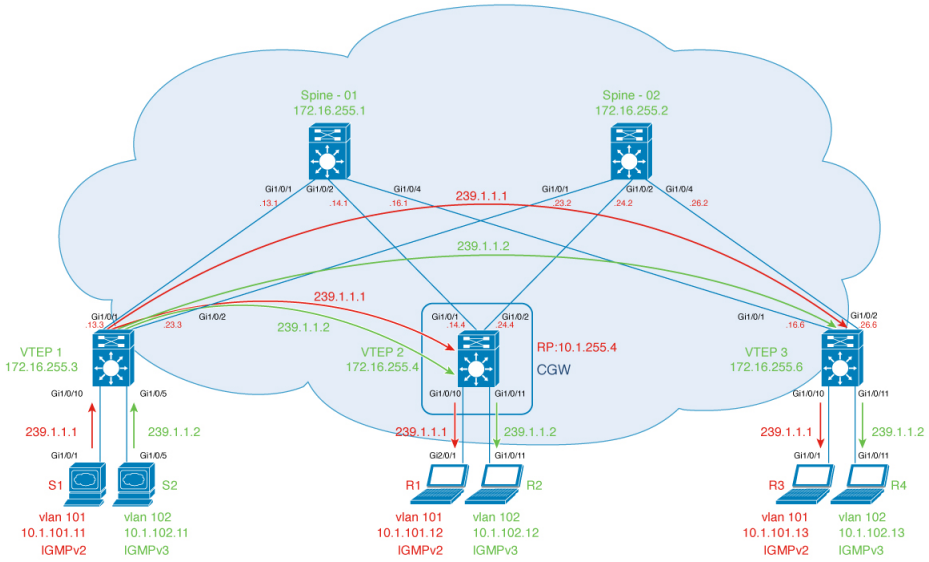
Configuration Examples for Optimized Layer 2 Overlay Multicast

The following sections provide examples to show how to configure Optimized Layer 2 Overlay Multicast with a Centralized Gateway, in a BGP EVPN VXLAN fabric.

Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication

This example shows how to configure Optimized Layer 2 Multicast across Layer 2 overlays in an EVPN VXLAN fabric that has ingress replication enabled in the underlay.

Figure 26: Topology for Optimized Layer 2 Overlay Multicast for IPv4 Traffic with Ingress Replication



The topology shows an EVPN VXLAN network with the source connected to Layer 2 VTEP 1 and a receiver connected to Layer 2 VTEP 3. Centralized Gateway is configured on VTEP 2. VLAN 101 has IGMPv2 and VLAN 102 has IGMPv3 enabled. The following tables provide sample configurations for the devices in this topology.

Table 27: Configure VTEP 1, CGW, and VTEP 3 for Optimized Layer 2 Multicast for IPv4, with Ingress Replication

VTEP 1	CGW	VTEP 3

VTEP 1	CGW	VTEP 3
Leaf-01#**show running-config** hostname Leaf-01 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.3 ip igmp snooping querier ! l2vpn evpn replication-type ingress multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 switchport access vlan 102 switchport mode access ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access	Leaf-02#**show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing vrf green ! l2vpn evpn replication-type ingress router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.1.255.4 255.255.255.255 ip pim sparse-mode ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0	Leaf-03#**show running-config** end hostname Leaf-03 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.6 ip igmp snooping querier ! l2vpn evpn replication-type ingress multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102

VTEP 1	CGW	VTEP 3
!	!	switchport mode access
interface nve1	interface GigabitEthernet1/0/2	!
no ip address	no switchport	interface nve1
source-interface Loopback1	ip address 172.16.24.4	no ip address
host-reachability protocol bgp	255.255.255.0	source-interface Loopback1
member vni 10101	ip ospf network point-to-point	host-reachability protocol bgp
ingress-replication	ip ospf 1 area 0	member vni 10101
member vni 10102	!	ingress-replication
ingress-replication	interface	member vni 10102
!	GigabitEthernet1/0/10	ingress-replication
router ospf 1	switchport access vlan 101	!
router-id 172.16.255.3	switchport mode access	router ospf 1
!	!	router-id 172.16.255.6
router bgp 65001	interface	!
bgp log-neighbor-changes	GigabitEthernet1/0/11	router bgp 65001
no bgp default ipv4-unicast	switchport access vlan 102	bgp log-neighbor-changes
neighbor 172.16.255.1	switchport mode access	no bgp default ipv4-unicast
remote-as 65001	spanning-tree portfast	neighbor 172.16.255.1
neighbor 172.16.255.1	!	remote-as 65001
update-source Loopback0	interface Vlan101	neighbor 172.16.255.1
neighbor 172.16.255.2	vrf forwarding green	update-source Loopback0
remote-as 65001	ip address 10.1.101.1	neighbor 172.16.255.2
neighbor 172.16.255.2	255.255.255.0	remote-as 65001
update-source Loopback0	ip pim sparse-mode	neighbor 172.16.255.2
!	!	update-source Loopback0
address-family ipv4	interface Vlan102	!
exit-address-family	vrf forwarding green	address-family ipv4
!	ip address 10.1.102.1	exit-address-family
address-family l2vpn evpn	255.255.255.0	!
neighbor 172.16.255.1 activate	ip pim sparse-mode	address-family l2vpn evpn
neighbor 172.16.255.1	ip igmp version 3	neighbor 172.16.255.1 activate
send-community both	!	neighbor 172.16.255.1
neighbor 172.16.255.2 activate	interface nve1	send-community both
neighbor 172.16.255.2	no ip address	neighbor 172.16.255.2 activate
send-community both	source-interface Loopback1	neighbor 172.16.255.2
exit-address-family	host-reachability protocol bgp	send-community both
!	member vni 10101	exit-address-family
end	ingress-replication	!
	member vni 10102	end
	ingress-replication	
	!	
	router ospf 1	
	router-id 172.16.255.4	
	!	
	router bgp 65001	
	bgp log-neighbor-changes	
	no bgp default ipv4-unicast	
	neighbor 172.16.255.1	
	remote-as 65001	
	neighbor 172.16.255.1	
	update-source Loopback0	
	neighbor 172.16.255.2	
	remote-as 65001	
	neighbor 172.16.255.2	
	update-source Loopback0	
	!	
	address-family ipv4	
	exit-address-family	
	!	
	address-family l2vpn evpn	
	neighbor 172.16.255.1 activate	
	neighbor 172.16.255.1	

VTEP 1	CGW	VTEP 3
	send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute static redistribute connected exit-address-family ! ip pim vrf green rp-address 10.1.255.4 ! end	

Table 28: Configure Spine Switch 1 and Spine Switch 2 for Optimized Layer 2 Multicast for IPv4, with Ingress Replication

Spine Switch 1	Spine Switch 2

Spine Switch 1	Spine Switch 2
Spine-01#**show running-config**	Spine-02#**show running-config**
hostname Spine-01	hostname Spine-02
!	!
ip routing	ip routing
!	!
interface Loopback0	interface Loopback0
ip address 172.16.255.1 255.255.255.255	ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface Loopback1	interface Loopback1
ip address 172.16.254.1 255.255.255.255	ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1
no switchport	no switchport
ip address 172.16.13.1 255.255.255.0	ip address 172.16.23.2 255.255.255.0
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2
no switchport	no switchport
ip address 172.16.14.1 255.255.255.0	ip address 172.16.24.2 255.255.255.0
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/4	interface GigabitEthernet1/0/4
no switchport	no switchport
ip address 172.16.16.1 255.255.255.0	ip address 172.16.26.2 255.255.255.0
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
router ospf 1	router ospf 1
router-id 172.16.255.1	router-id 172.16.255.2
!	!
router bgp 65001	router bgp 65001
bgp router-id 172.16.255.1	bgp router-id 172.16.255.2
bgp log-neighbor-changes	bgp log-neighbor-changes
no bgp default ipv4-unicast	no bgp default ipv4-unicast
neighbor 172.16.255.2 remote-as 65001	neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0	neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.3 remote-as 65001	neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0	neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001	neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0	neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001	neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0	neighbor 172.16.255.6 update-source Loopback0
!	!
address-family ipv4	address-family ipv4
exit-address-family	exit-address-family
!	!
address-family l2vpn evpn	address-family l2vpn evpn
neighbor 172.16.255.2 activate	neighbor 172.16.255.1 activate
neighbor 172.16.255.2 send-community both	neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 route-reflector-client	neighbor 172.16.255.1 route-reflector-client
neighbor 172.16.255.3 activate	neighbor 172.16.255.3 activate
neighbor 172.16.255.3 send-community both	neighbor 172.16.255.3 send-community both
neighbor 172.16.255.3 route-reflector-client	neighbor 172.16.255.3 route-reflector-client
neighbor 172.16.255.4 activate	neighbor 172.16.255.4 activate
neighbor 172.16.255.4 send-community both	neighbor 172.16.255.4 send-community both
neighbor 172.16.255.4 route-reflector-client	neighbor 172.16.255.4 route-reflector-client
neighbor 172.16.255.6 activate	neighbor 172.16.255.6 activate
neighbor 172.16.255.6 send-community both	neighbor 172.16.255.6 send-community both

Spine Switch 1	Spine Switch 2
neighbor 172.16.255.6 route-reflector-client exit-address-family ! end	neighbor 172.16.255.6 route-reflector-client exit-address-family ! end

To return, click Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication, on page 215

To return to the Configuration Examples section, click Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214.

Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay

The following sections provide sample output of the **show** commands to verify Optimized Layer 2 Overlay Multicast with Ingress Replication on the devices in the topology configured above.

Outputs to verify configuration on VTEP 1

Outputs to verify configuration on CGW

Outputs to verify configuration on VTEP 3

Outputs to Verify Configuration on VTEP 1

```
Leaf-01# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4


Leaf-01# show ip igmp snooping querier vlan 101
IP address               : 172.16.254.3
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2


Leaf-01# show ip igmp snooping groups vlan 101
Vlan      Group                   Type        Version     Port List
-------------------------------------------------------------------------
101       239.1.1.1               igmp        v2          Tu0


Leaf-01# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
```

```
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 0 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 1 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 0 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-01# show l2vpn evpn multicast local address 239.1.1.1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
```

```
Leaf-01# show l2vpn evpn multicast remote address 239.1.1.1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         IGMPv2    N/A     (*, 239.1.1.1)
101   101   172.16.254.6         IGMPv2    N/A     (*, 239.1.1.1)
```

```
Leaf-01# show l2route evpn multicast routes group 239.1.1.1
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
101   0          239.1.1.1       *               V:10101 172.16.254.6, V:10101 172.16.254.4
```

```
Leaf-01# show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin              Group           Filter     Source(s)
----- ---------- ------------------- --------------- ---------- --------------------
101   0          172.16.254.4        239.1.1.1       N/A        (*)IGMPv2
101   0          172.16.254.6        239.1.1.1       N/A        (*)IGMPv2
```

```
Leaf-01# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 87
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:02:18 UTC
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 77
```

```
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 4 2022 20:02:18 UTC
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:02:18 UTC


Leaf-01# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 192
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:03:54 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 191
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:03:54 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 4 2022 20:03:54 UTC
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay, on page 222.

Outputs to Verify Configuration on CGW

```
Leaf-02# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d Vl102


Leaf-02# show ip igmp snooping querier vlan 102
IP address               : 10.1.102.1
IGMP version             : v3
Port                     : Router
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2


Leaf-02# show ip igmp snooping groups vlan 102
Vlan      Group                    Type        Version     Port List
------------------------------------------------------------------------------
102       239.1.1.2                igmp        v3          Gi1/0/1, Tu0

Leaf-02#show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.4:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan102
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 0 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 0 MAC/IP, 1 IMET, 0 EAD


Leaf-02# show l2vpn evpn multicast local address 239.1.1.2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
```

```
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11        IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)


Leaf-02# show l2vpn evpn multicast remote address 239.1.1.2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.6         IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)


Leaf-02# show l2route evpn multicast routes group 239.1.1.2
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
102   0          239.1.1.2       *               Gi1/0/11:102, V:10102 172.16.254.6


Leaf-02# show l2route evpn multicast smet group 239.1.1.2
EVI   ETAG       Origin               Group           Filter     Source(s)
----- ---------- -------------------- --------------- ---------- --------------------
102   0          Gi1/0/11:102         239.1.1.2       INCLUDE    10.1.102.11
102   0          172.16.254.6         239.1.1.2       INCLUDE    10.1.102.11


Leaf-02# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 116
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:02:21 UTC


Leaf-02# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.6]/27, version 186
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.6:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.6]/27 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:03:50 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.6]/27, version 184
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
```

```
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 4 2022 20:03:50 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:03:50 UTC


Leaf-02# show ip mroute vrf green 239.1.1.2
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.2), 01:05:00/00:02:58, RP 10.1.255.4, flags: SJC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan102, Forward/Sparse, 00:44:20/00:02:58, flags:
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay, on page 222.

Outputs to Verify Configuration on VTEP 3

```
Leaf-03# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                  EVI   VLAN  MAC Address    Source
----- ---------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                               102   102   7c21.0dbd.954d 172.16.254.4

Leaf-03#show ip igmp snooping querier vlan 102
IP address               : 172.16.254.6
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2
```

```
Leaf-03# show ip igmp snooping groups vlan 102
Vlan      Group                Type         Version     Port List
------------------------------------------------------------------------
102       239.1.1.2            igmp         v3          Gi1/0/11, Tu0


Leaf-03# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.6:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 0 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 1 MAC/IP, 1 IMET, 0 EAD


Leaf-03# show l2vpn evpn multicast local address 239.1.1.2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11        IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)


Leaf-03#show l2vpn evpn multicast remote address 239.1.1.2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.4         IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)


Leaf-03#show l2route evpn multicast routes group 239.1.1.2
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
102   0          239.1.1.2       *               Gi1/0/11:102, V:10102 172.16.254.4


Leaf-03# show l2route evpn multicast smet group 239.1.1.2
EVI   ETAG       Origin               Group           Filter     Source(s)
----- ---------- -------------------- --------------- ---------- --------------------
```

```
102   0          Gi1/0/11:102         239.1.1.2       INCLUDE    10.1.102.11
102   0          172.16.254.4         239.1.1.2       INCLUDE    10.1.102.11
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 208
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:02:21 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 4 2022 20:02:21 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 210
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:02:21 UTC
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.6]/27, version 230
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 4 2022 20:03:50 UTC
```

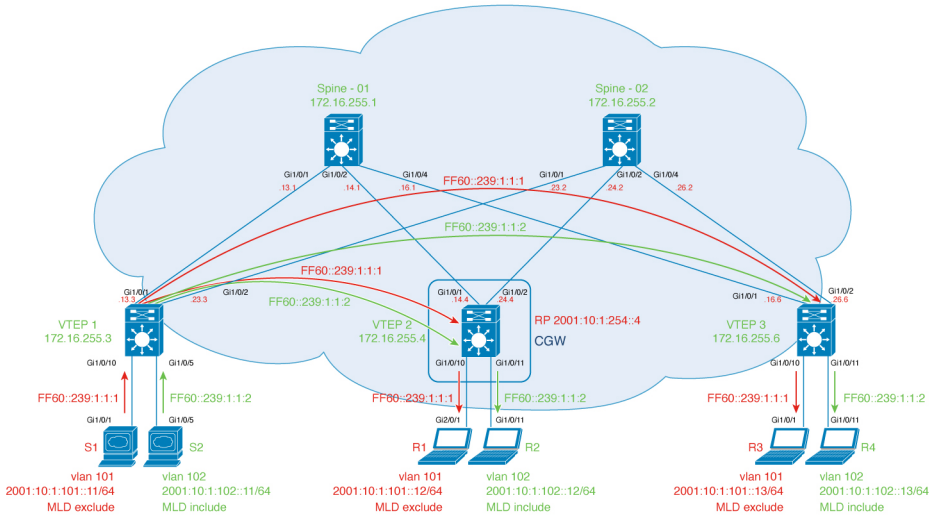
To return, click Configuration Example for Optimized Layer 2 Multicast with Ingress Replication in the Underlay.

Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6

This example shows how to configure Optimized Layer 2 Overlay Multicast across the Layer 2 overlays in an EVPN VXLAN fabric that has ingress replication enabled in the underlay. The example configuration is for both IPv4 and IPv6 multicast traffic.

For IPv4 multicast traffic, refer Topology for Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4.

Figure 27: Topology for Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv6 Traffic



The topology shows an EVPN VXLAN network with the source connected to Layer 2 VTEP 1 and receivers connected to Layer 2 VTEP 3. A Centralized Gateway is configured on VTEP 2. VLAN 101 has MLD exclude mode and VLAN 102 has MLD include mode. The following tables provide sample configurations for IPv4 and IPv6 multicast traffic for the devices in this topology:

Table 29: Configure VTEP 1, CGW, and VTEP 3 for Optimized Layer 2 Multicast for IPv4 and IPv6, with Ingress Replication

VTEP1	CGW	VTEP 3

VTEP1	CGW	VTEP 3
Leaf-01#**show running-config** hostname Leaf-01 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.3 ip igmp snooping querier ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ! l2vpn evpn replication-type ingress router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 !	Leaf-02#**show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ipv6 unicast-routing ipv6 mld snooping ipv6 multicast-routing vrf green ! l2vpn evpn replication-type ingress router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 !	Leaf-03#**show running-config** hostname Leaf-03 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.6 ip igmp snooping querier ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ! l2vpn evpn replication-type ingress router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 !

VTEP1	CGW	VTEP 3
interface GigabitEthernet1/0/5 switchport access vlan 102 switchport mode access ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end	interface Loopback255 vrf forwarding green ip address 10.1.255.4 255.255.255.255 ip pim sparse-mode ipv6 address 2001:10:1:255::4/128 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address 2001:10:1:101::1/64 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ipv6 address 2001:10:1:102::1/64 ! interface Vlan901 no ip address ipv6 enable ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication	interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end

VTEP1	CGW	VTEP 3
	! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute static redistribute connected exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim vrf green rp-address 10.1.255.4 ipv6 pim vrf green rp-address 2001:10:1:255::4 ! end	

Table 30: Configure Spine Switch 1 and Spine Switch 2 for Optimized Layer 2 Multicast for IPv4 and IPv6, with Ingress Replication

Spine Switch 1	Spine Switch 2

Spine Switch 1	Spine Switch 2
Spine-01#**show running-config**	Spine-02#**show running-config**
hostname Spine-01	hostname Spine-02
!	!
ip routing	ip routing
!	!
interface Loopback0	interface Loopback0
ip address 172.16.255.1 255.255.255.255	ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface Loopback1	interface Loopback1
ip address 172.16.254.1 255.255.255.255	ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1
no switchport	no switchport
ip address 172.16.13.1 255.255.255.0	ip address 172.16.23.2 255.255.255.0
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2
no switchport	no switchport
ip address 172.16.14.1 255.255.255.0	ip address 172.16.24.2 255.255.255.0
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/4	interface GigabitEthernet1/0/4
no switchport	no switchport
ip address 172.16.16.1 255.255.255.0	ip address 172.16.26.2 255.255.255.0
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
router ospf 1	router ospf 1
router-id 172.16.255.1	router-id 172.16.255.2
!	!
router bgp 65001	router bgp 65001
bgp router-id 172.16.255.1	bgp router-id 172.16.255.2
bgp log-neighbor-changes	bgp log-neighbor-changes
no bgp default ipv4-unicast	no bgp default ipv4-unicast
neighbor 172.16.255.2 remote-as 65001	neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0	neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.3 remote-as 65001	neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0	neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001	neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0	neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001	neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0	neighbor 172.16.255.6 update-source Loopback0
!	!
address-family ipv4	address-family ipv4
exit-address-family	exit-address-family
!	!
address-family l2vpn evpn	address-family l2vpn evpn
neighbor 172.16.255.2 activate	neighbor 172.16.255.1 activate
neighbor 172.16.255.2 send-community both	neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 route-reflector-client	neighbor 172.16.255.1 route-reflector-client
neighbor 172.16.255.3 activate	neighbor 172.16.255.3 activate
neighbor 172.16.255.3 send-community both	neighbor 172.16.255.3 send-community both
neighbor 172.16.255.3 route-reflector-client	neighbor 172.16.255.3 route-reflector-client
neighbor 172.16.255.4 activate	neighbor 172.16.255.4 activate
neighbor 172.16.255.4 send-community both	neighbor 172.16.255.4 send-community both
neighbor 172.16.255.4 route-reflector-client	neighbor 172.16.255.4 route-reflector-client
neighbor 172.16.255.6 activate	neighbor 172.16.255.6 activate
neighbor 172.16.255.6 send-community both	neighbor 172.16.255.6 send-community both

Spine Switch 1	Spine Switch 2
neighbor 172.16.255.6 route-reflector-client exit-address-family ! end	neighbor 172.16.255.6 route-reflector-client exit-address-family ! end

To return, click Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6, on page 230.

To return to the Configuration Examples section, click Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214.

Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic

The following sections provide sample output of **show** commands to verify Optimized Layer 2 Overlay Multicast with Ingress Replication on the devices in the topology configured above.

To see the output of the **show** commands for IPv4 multicast traffic, go to Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay, on page 222.

The output of the **show** commands for IPv6 multicast traffic are provided in the following sections.

Outputs to verify configuration on VTEP 1

Outputs to verify configuration on CGW

Outputs to verify configuration on VTEP 3

Outputs to verify configuration on Spine Switch 1

Outputs to verify configuration on Spine Switch 2

Outputs to Verify Configuration on VTEP 1

Verify the configuration with MLD exclude.

```
Leaf-01# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 172.16.254.4

Leaf-01# show ipv6 mld snooping querier vlan 101
IP address               : FE80:0:4A56:0:12B3:D5FF:FE6A:8F80
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2

Leaf-01# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
-----------------------------------------
Total number of channels: 2
Total number of hosts   : 2

Source/Group                    Interface Reporter                      Vlan Uptime
Last-Join/
```

```
                                                                                                  Last-Leave
-------------------------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Tu0       FE80::AC10:FE04                 101  20:14:25
20:14:25 /
                                                                                                      -


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE06                 101  20:13:50
20:13:50 /
                                                                                                      -


::/FF06:239:1:1::1              Tu0       FE80::AC10:FE04                 101  00:00:48
20:15:23 /
                                                                                                20:14:34


::/FF06:239:1:1::1              Tu0       FE80::AC10:FE06                 101  00:01:11
20:15:10 /
                                                                                                20:13:59


Leaf-01# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD

Leaf-01# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
```

```
Leaf-01# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         MLDv2     EXCLUDE (*, FF06::239:1:1:1)
101   101   172.16.254.6         MLDv2     EXCLUDE (*, FF06::239:1:1:1)


Leaf-01# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
-------------------------------------- --------------------
101   0          FF06::239:1:1:1                          ::
       V:10101 172.16.254.4, Gi1/0/10:101, V:10101 172.16.254.6


Leaf-01# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin               Group                                    Filter
Source(s)
----- ---------- -------------------- ---------------------------------------- ----------
--------------------
101   0          172.16.254.4         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
101   0          172.16.254.6         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2


Leaf-01# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 121
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:06:25 UTC
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 120
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:06:25 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 6 2022 16:06:25 UTC
```

```
Leaf-01# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 131
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:07:00 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 129
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 6 2022 16:07:00 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:07:00 UTC
```

Verify the configuration with MLD include.

```
Leaf-01# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d 172.16.254.4
  Y   2001:10:1:102::1                        102   102   7c21.0dbd.954d 172.16.254.4

Leaf-01# show ipv6 mld snooping querier vlan 102
IP address               : FE80:0:4A56:0:12B3:D5FF:FE6A:8F80
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2

Leaf-01# show ipv6 mld snooping membership vlan 102 source 2001:10:1:102::11 group
FF06::239:1:1:2
Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------
```

```
2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE04                 102  00:02:56
00:02:56 /
                                                                                       -


2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE06                 102  00:02:56
00:02:56 /

Leaf-01# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.3:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    Pseudoports:
      GigabitEthernet1/0/5 service instance 102
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD

Leaf-01#show l2vpn evpn multicast local address FF06::239:1:1:2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------

Leaf-01# show l2vpn evpn multicast remote address FF06::239:1:1:2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.4         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
102   102   172.16.254.6         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)

Leaf-01# show l2route evpn multicast routes group FF06::239:1:1:2
EVI   ETAG       Group                                   Source
       Next-hop(s)
----- ---------- ----------------------------------------
---------------------------------------- --------------------
102   0          FF06::239:1:1:2                         ::
```

```
       V:10102 172.16.254.4, V:10102 172.16.254.6, Gi1/0/5:102

Leaf-01# show l2route evpn multicast smet group FF06::239:1:1:2EVI   ETAG       Origin
          Group                                  Filter     Source(s)
----- ---------- -------------------- ------------------------------------- ----------
--------------------
102   0          172.16.254.4         FF06::239:1:1:2                       INCLUDE
2001:10:1:102::11
102   0          172.16.254.6         FF06::239:1:1:2                       INCLUDE
2001:10:1:102::11


Leaf-01# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 87
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:45:47 UTC
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 35
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 13:44:53 UTC
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:44:53 UTC


Leaf-01# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 95
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51
 (global)
```

```
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:45:47 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 39
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 13:44:53 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:44:53 UTC
```

To return click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 237.

Outputs to Verify Configuration on CGW

Verify the configuration with MLD exclude.

```
Leaf-02# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 Vl101
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 Vl101


Leaf-02# show ipv6 mld snooping querier vlan 101
IP address               : FE80::46D3:CAFF:FE28:6CC1
MLD  version             : v2
Port                     : Gi1/0/10
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-02# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
----------------------------------------
Total number of channels: 2
Total number of hosts   : 2

Source/Group                    Interface Reporter                      Vlan Uptime
```

```
Last-Join/
                                                                                                  Last-Leave
-------------------------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Gi1/0/10  FE80::46D3:CAFF:FE28:6CC1       101  00:00:00
00:00:05 /
                                                                                                  00:00:05


::/FF06:239:1:1::1              Gi1/0/10  FE80::46D3:CAFF:FE28:6CC1       101  00:00:47
20:16:44 /
                                                                                                  20:15:56


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE06                 101  20:15:12
20:15:12 /
                                                                                                  -


::/FF06:239:1:1::1              Tu0       FE80::AC10:FE06                 101  00:00:00
20:15:21 /
                                                                                                  20:15:21


Leaf-02# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Enabled (Asymmetric)
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-02# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        MLDv2     EXCLUDE (*, FF06::239:1:1:1)

Leaf-02# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.6         MLDv2     EXCLUDE (*, FF06::239:1:1:1)

Leaf-02# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
--------------------------------------- --------------------
101   0          FF06::239:1:1:1                          ::
       Gi1/0/10:101, V:10101 172.16.254.3, V:10101 172.16.254.6

Leaf-02# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin               Group                                    Filter
Source(s)
----- ---------- -------------------- ---------------------------------------- ----------
--------------------
101   0          Gi1/0/10:101         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
101   0          172.16.254.6         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2

Leaf-02# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 117
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:06:25 UTC

Leaf-02# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 126
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
```

```
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:07:00 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 124
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 6 2022 16:07:00 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:07:00 UTC


Leaf-02# show ipv6 mroute vrf green FF06::239:1:1:1
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06::239:1:1:1), 20:17:15/never, RP 2001:10:1:255::4, flags: SCJ
  Incoming interface: Tunnel6
  RPF nbr: 2001:10:1:255::4
  Immediate Outgoing interface list:
    Vlan101, Forward, 20:17:15/never

(2001:10:1:101::11, FF06::239:1:1:1), 00:00:33/00:02:56, flags: SFJT
  Incoming interface: Vlan101
  RPF nbr: 2001:10:1:101::11
  Outgoing interface list: Null
```

Configuration with MLD include

```
Leaf-02# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d Vl102

Leaf-02# show ipv6 mld snooping querier vlan 102
IP address               : FE80::46D3:CAFF:FE28:6CC2
```

```
MLD  version             : v2
Port                     : Gi1/0/11
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2
```

```
Leaf-02# show ipv6 mld snooping membership vlan 102 source 2001:10:1:102::11 group
FF06::239:1:1:2
Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------

2001:10:1:102::11/FF06::239:1:1:2 \
                                Gi1/0/11  FE80::46D3:CAFF:FE28:6CC2       102  00:10:26
00:00:09 /
                                                                                        -


2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE06                 102  00:08:42
00:08:42 /
                                                                                        -
```

```
Leaf-02# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.4:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan102
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Enabled (Asymmetric)
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-02# show l2vpn evpn multicast local address FF06::239:1:1:2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11       MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)


Leaf-02# show l2vpn evpn multicast remote address FF06::239:1:1:2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.6         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)


Leaf-02# show l2route evpn multicast routes group FF06::239:1:1:2
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
--------------------------------------- --------------------
102   0          FF06::239:1:1:2                          ::
       Gi1/0/11:102, V:10102 172.16.254.3, V:10102 172.16.254.6


Leaf-02# show l2route evpn multicast smet group FF06::239:1:1:2
EVI   ETAG       Origin               Group                                    Filter
Source(s)
----- ---------- -------------------- ---------------------------------------- ----------
--------------------
102   0          Gi1/0/11:102         FF06::239:1:1:2                          INCLUDE
2001:10:1:102::11
102   0          172.16.254.6         FF06::239:1:1:2                          INCLUDE
2001:10:1:102::11

Leaf-02# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 143
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:39:42 UTC


Leaf-02# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 164
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
```

```
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:41:25 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 160
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 13:41:25 UTC
  Refresh Epoch 3
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:44:52 UTC


Leaf-02# show ipv6 mroute vrf green FF06::239:1:1:2
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(2001:10:1:102::11, FF06::239:1:1:2), 00:14:17/00:00:37, flags: SFTI
  Incoming interface: Vlan102
  RPF nbr: 2001:10:1:102::11
  Outgoing interface list: Null
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 237.

Outputs to Verify Configuration on VTEP 3

MLD exclude

```
Leaf-03# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 172.16.254.4
```

```
Leaf-03# show ipv6 mld snooping querier vlan 101
IP address               : FE80:0:8155:0:E75:BDFF:FE67:EF00
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-03# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
-----------------------------------------
Total number of channels: 2
Total number of hosts   : 2

Source/Group                    Interface Reporter                      Vlan Uptime
Last-Join/
                                                                              Last-Leave
-------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Gi1/0/10  FE80::EEE1:A9FF:FE37:92C1     101  00:00:00
00:00:57 /
                                                                              00:00:57


::/FF06:239:1:1::1              Gi1/0/10  FE80::EEE1:A9FF:FE37:92C1     101  00:01:07
20:25:57 /
                                                                              20:24:49


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE04               101  20:25:16
20:25:16 /
                                                                                    -


::/FF06:239:1:1::1              Tu0       FE80::AC10:FE04               101  00:00:48
20:26:13 /
                                                                              20:25:25


Leaf-03# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.6:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
```

```
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD

Leaf-03# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        MLDv2     EXCLUDE (*, FF06::239:1:1:1)

Leaf-03# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         MLDv2     EXCLUDE (*, FF06::239:1:1:1)

Leaf-03# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
--------------------------------------- --------------------
101   0          FF06::239:1:1:1                          ::
       V:10101 172.16.254.4, V:10101 172.16.254.3, Gi1/0/10:101

Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin               Group                                    Filter
Source(s)
----- ---------- -------------------- ---------------------------------------- ----------
--------------------
101   0          Gi1/0/10:101         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
101   0          172.16.254.4         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2

Leaf-03# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 105
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:06:25 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
```

```
      Updated on Apr 6 2022 16:06:25 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 106
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:06:25 UTC


Leaf-03# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 111
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 6 2022 16:07:00 UTC
```

Verify the configuration with MLD include

```
Leaf-03# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d 172.16.254.4
  Y   2001:10:1:102::1                        102   102   7c21.0dbd.954d 172.16.254.4


Leaf-03# show ipv6 mld snooping querier vlan 102
IP address               : FE80:0:8155:0:E75:BDFF:FE67:EF00
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-03# show ipv6 mld snooping membership vlan 102 source 2001:10:1:102::11 group
FF06::239:1:1:2
Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                 Last-Leave
-------------------------------------------------------------------------------------------
```

```
2001:10:1:102::11/FF06::239:1:1:2 \
                                Gi1/0/11 FE80::EEE1:A9FF:FE37:92C2       102  00:11:00
00:00:25 /
                                                                                      -


2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE04                 102  00:12:44
00:12:44 /
                                                                                      -


Leaf-03# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.6:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD

Leaf-03# show l2vpn evpn multicast local address FF06::239:1:1:2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11       MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)

Leaf-03# show l2vpn evpn multicast remote address FF06::239:1:1:2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.4         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)

Leaf-03# show l2route evpn multicast routes group FF06::239:1:1:2
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
--------------------------------------- --------------------
```

```
102   0          FF06::239:1:1:2                          ::
       V:10102 172.16.254.4, V:10102 172.16.254.3, Gi1/0/11:102


Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:2
EVI   ETAG       Origin               Group                                    Filter
Source(s)
----- ---------- -------------------- ---------------------------------------- ----------
--------------------
102   0          Gi1/0/11:102         FF06::239:1:1:2                          INCLUDE
2001:10:1:102::11
102   0          172.16.254.4         FF06::239:1:1:2                          INCLUDE
2001:10:1:102::11


Leaf-03# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 138
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 13:39:42 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:39:42 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 140
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.254.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 13:39:42 UTC


Leaf-03# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 146
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
```

```
Local
  :: (via default) from 0.0.0.0 (172.16.255.6)
    Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
    IGMP/MLD v2
    Extended Community: RT:65001:102 ENCAP:8
    Local irb vxlan vtep:
      vrf:not found, l3-vni:0
      local router mac:0000.0000.0000
      core-irb interface:(not found)
      vtep-ip:172.16.254.6
    rx pathid: 0, tx pathid: 0x0
    Updated on Apr 7 2022 13:41:25 UTC
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 237.

Outputs to Verify Configuration on Spine Switch 1

Verify the configuration with MLD exclude.

```
Spine-01# show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 79, main routing table version 79
44 network entries using 16896 bytes of memory
88 path entries using 19712 bytes of memory
11/11 BGP path/bestpath attribute entries using 3256 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
7 BGP extended community entries using 264 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 40248 total bytes of memory
BGP activity 54/10 prefixes, 108/20 paths, scan interval 60 secs
47 networks peaked at 16:07:10 Apr 6 2022 UTC (20:25:33.334 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001    1423    1428       79    0    0 20:43:17        44
172.16.255.3    4        65001    1384    1431       79    0    0 20:46:03        10
172.16.255.4    4        65001    1400    1449       79    0    0 20:46:02        22
172.16.255.6    4        65001    1389    1428       79    0    0 20:45:54        12
```

Verify the configuration with MLD include.

```
Spine-01#sh bgp l2vpn evpn all summ
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 79, main routing table version 79
44 network entries using 16896 bytes of memory
88 path entries using 19712 bytes of memory
11/11 BGP path/bestpath attribute entries using 3256 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
7 BGP extended community entries using 264 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 40248 total bytes of memory
BGP activity 54/10 prefixes, 108/20 paths, scan interval 60 secs
47 networks peaked at 16:07:10 Apr 6 2022 UTC (21:19:08.411 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001    1482    1487       79    0    0 21:36:52        44
```

```
172.16.255.3    4        65001    1443    1491       79    0    0 21:39:38        10
172.16.255.4    4        65001    1458    1508       79    0    0 21:39:37        22
172.16.255.6    4        65001    1447    1487       79    0    0 21:39:30        12
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 237.

Outputs to Verify Configuration on Spine Switch 2

Verify the configuration with MLD exclude.

```
Spine-02# show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 79, main routing table version 79
44 network entries using 16896 bytes of memory
88 path entries using 19712 bytes of memory
11/11 BGP path/bestpath attribute entries using 3256 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
7 BGP extended community entries using 264 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP usinory
BGP activity 54/10 prefixes, 108/20 paths, scan interval 60 secs
47 networkg 40248 total bytes of mems peaked at 16:07:00 Apr 6 2022 UTC (20:26:17.324 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1429    1424       79    0    0 20:43:51        44
172.16.255.3    4        65001    1384    1429       79    0    0 20:43:47        10
172.16.255.4    4        65001    1401    1446       79    0    0 20:43:45        22
172.16.255.6    4        65001    1385    1427       79    0    0 20:43:48        12
```

Verify the configuration with MLD include.

```
Spine-02#show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 79, main routing table version 79
44 network entries using 16896 bytes of memory
88 path entries using 19712 bytes of memory
11/11 BGP path/bestpath attribute entries using 3256 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
7 BGP extended community entries using 264 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 40248 total bytes of memory
BGP activity 54/10 prefixes, 108/20 paths, scan interval 60 secs
47 networks peaked at 16:07:00 Apr 6 2022 UTC (21:19:32.246 ago)

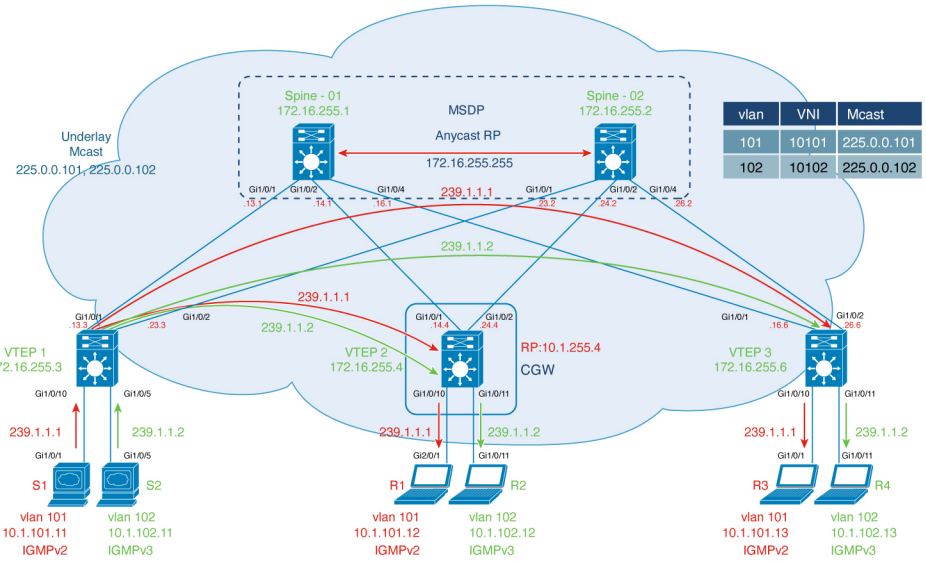
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1487    1482       79    0    0 21:37:06        44
172.16.255.3    4        65001    1443    1489       79    0    0 21:37:02        10
172.16.255.4    4        65001    1459    1504       79    0    0 21:37:00        22
172.16.255.6    4        65001    1444    1485       79    0    0 21:37:03        12
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 237.

Example 3: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication

This example shows how to configure Optimized Layer 2 Multicast for IPv4 multicast traffic across the Layer 2 overlays in an EVPN VXLAN fabric that has underlay multicast replication enabled.

Figure 28: Topology for Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication



The topology shows an EVPN VXLAN network with the source connected to Layer 2 VTEP 1 and a receiver connected to Layer 2 VTEP 3. Centralized Gateway is configured on VTEP 2. The multicast group for VLAN 101 is 225.0.0.101 and for VLAN 102 is 225.0.0.102. The following tables provide sample configurations for the devices in this topology:

Table 31: Configure VTEP 1, CGW, and VTEP 3 for Optimized Layer 2 Overlay Multicast for IPv4 Traffic with Underlay Multicast Replication

VTEP 1	CGW	VTEP 3

VTEP 1	CGW	VTEP 3
Leaf-01#**show running-config** hostname Leaf-01 ! ip routing ! ip multicast-routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.3 ip igmp snooping querier ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5	Leaf-02#**show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.1.255.4 255.255.255.255 ip pim sparse-mode ! interface GigabitEthernet1/0/1 no switchport	Leaf-03#**show running-config** hostname Leaf-03 ! ip routing ! ip multicast-routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.6 ip igmp snooping querier ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface

VTEP 1	CGW	VTEP 3
switchport access vlan 102 switchport mode access ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-address 172.16.255.255 ! end	ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1	GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-address 172.16.255.255 ! end

VTEP 1	CGW	VTEP 3
	update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ip pim vrf green rp-address 10.1.255.4 ! end	

Table 32: Configure Spine Switch 1 and Spine Switch 2 for Optimized Layer 2 Overlay Multicast for IPv4 Traffic, with Underlay Multicast Replication

Spine Switch 1	Spine Switch 2

Spine Switch 1	Spine Switch 2
Spine-01#**show running-config**	Spine-02#**show running-config**
hostname Spine-01	hostname Spine-02
!	!
ip routing	ip routing
!	!
ip multicast-routing	ip multicast-routing
!	!
interface Loopback0	interface Loopback0
ip address 172.16.255.1 255.255.255.255	ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface Loopback1	interface Loopback1
ip address 172.16.254.1 255.255.255.255	ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface Loopback2	interface Loopback2
ip address 172.16.255.255 255.255.255.255	ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode	ip pim sparse-mode
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1
no switchport	no switchport
ip address 172.16.13.1 255.255.255.0	ip address 172.16.23.2 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2
no switchport	no switchport
ip address 172.16.14.1 255.255.255.0	ip address 172.16.24.2 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/4	interface GigabitEthernet1/0/4
no switchport	no switchport
ip address 172.16.16.1 255.255.255.0	ip address 172.16.26.2 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
router ospf 1	router ospf 1
router-id 172.16.255.1	router-id 172.16.255.2
!	!
router bgp 65001	router bgp 65001
bgp router-id 172.16.255.1	bgp router-id 172.16.255.2
bgp log-neighbor-changes	bgp log-neighbor-changes
no bgp default ipv4-unicast	no bgp default ipv4-unicast
neighbor 172.16.255.2 remote-as 65001	neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0	neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.3 remote-as 65001	neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0	neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001	neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0	neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001	neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0	neighbor 172.16.255.6 update-source Loopback0
!	!
address-family ipv4	address-family ipv4
exit-address-family	exit-address-family
!	!
address-family l2vpn evpn	address-family l2vpn evpn
neighbor 172.16.255.2 activate	neighbor 172.16.255.1 activate

Spine Switch 1	Spine Switch 2
neighbor 172.16.255.2 send-community both neighbor 172.16.255.2 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end	neighbor 172.16.255.1 send-community both neighbor 172.16.255.1 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end

To return, click Configuration Example for Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication

To return to the Configuration Examples section, click Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214.

Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay

The following sections provide sample output of the **show** commands to verify the configuration of Optimized Layer 2 Overlay Multicast with default MDT on the devices in the topology configured above.

Outputs to verify configuration on VTEP 1

Outputs to verify configuration on CGW

Outputs to verify configuration on VTEP 3

Outputs to Verify Configuration on VTEP 1

```
Leaf-01#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.13.1       GigabitEthernet1/0/1     00:16:02/00:01:28 v2    1 / S P G
172.16.23.2       GigabitEthernet1/0/2     00:16:02/00:01:26 v2    1 / S P G


Leaf-01#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-01#show ip mroute 225.0.0.101
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 00:15:58/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:15:58/00:02:03, flags:

(172.16.254.3, 225.0.0.101), 00:11:19/00:02:10, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:11:19/00:03:01, flags:


Leaf-01#show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                  EVI   VLAN  MAC Address    Source
----- ---------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                               101   101   7c21.0dbd.9541 172.16.254.4

Leaf-01#show ip igmp snooping querier vlan 101
IP address               : 172.16.254.3
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2

Leaf-01#show ip igmp snooping groups vlan 101
Vlan      Group                    Type        Version     Port List
------------------------------------------------------------------------------
101       239.1.1.1                igmp        v2          Tu0


Leaf-01#show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
```

```
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    MCAST IP:        225.0.0.101
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 0 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 1 MAC/IP, 0 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 0 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-01#show l2vpn evpn multicast remote address 239.1.1.1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         IGMPv2    N/A     (*, 239.1.1.1)
101   101   172.16.254.6         IGMPv2    N/A     (*, 239.1.1.1)
```

```
Leaf-01# show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin               Group           Filter     Source(s)
----- ---------- -------------------- --------------- ---------- --------------------
101   0          172.16.254.4         239.1.1.1       N/A        (*)IGMPv2
101   0          172.16.254.6         239.1.1.1       N/A        (*)IGMPv2
```

```
Leaf-01#show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 31
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 19:57:09 UTC
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 11
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 5 2022 19:57:09 UTC
```

```
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 19:57:09 UTC


Leaf-01#show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 65
Paths: (1 available, best #1, table evi_101)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:07:46 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 57
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 5 2022 20:07:46 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:07:46 UTC
```

To return, click Verifying Optimized Layer 2 Multicast with Default MDT in the Underlay .

Outputs to Verify Configuration on CGW

```
Leaf-02#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
```

```
172.16.14.1       GigabitEthernet1/0/1     00:21:12/00:01:43 v2    1 / S P G
172.16.24.2       GigabitEthernet1/0/2     00:21:11/00:01:40 v2    1 / S P G


Leaf-02#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-02#show ip mroute 225.0.0.101
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 00:21:16/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:21:16/00:02:45, flags:

(172.16.254.3, 225.0.0.101), 00:15:36/00:01:15, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:15:36/00:02:23, flags:


Leaf-02#show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 Vl101

Leaf-02#show ip igmp snooping querier vlan 101
IP address               : 10.1.101.1
IGMP version             : v2
Port                     : Router
Max response time        : 10s


Leaf-02#show ip igmp snooping groups vlan 101
Vlan      Group                   Type        Version     Port List
------------------------------------------------------------------------------
101       239.1.1.1               igmp        v2          Gi1/0/10, Tu0


Leaf-02#show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
```

```
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:         Vlan901
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            7c21.0dbd.9548
    Core Vlan:       901
    L2 VNI:          10101
    L3 VNI:          50901
    VTEP IP:         172.16.254.4
    MCAST IP:        225.0.0.101
    VRF:             green
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 0 MAC/IP, 0 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 0 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-02#show l2vpn evpn multicast local address 239.1.1.1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        IGMPv2    N/A     (*, 239.1.1.1)

Leaf-02#show l2vpn evpn multicast remote address 239.1.1.1EVI   VLAN  Originator
 Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.6         IGMPv2    N/A     (*, 239.1.1.1)


Leaf-02#show l2route evpn multicast routes group 239.1.1.1
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
101   0          239.1.1.1       *               Gi1/0/10:101, V:10101 225.0.0.101


Leaf-02#show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin              Group           Filter     Source(s)
----- ---------- ------------------- --------------- ---------- --------------------
101   0          Gi1/0/10:101        239.1.1.1       N/A        (*)IGMPv2
101   0          172.16.254.6        239.1.1.1       N/A        (*)IGMPv2


Leaf-02#show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 12
```

```
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:green, l3-vni:50901
        local router mac:7C21.0DBD.9548
        core-irb interface:Vlan901
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 19:56:23 UTC


Leaf-02#show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 62
Paths: (1 available, best #1, table evi_101)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:07:46 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 54
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 5 2022 20:07:46 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:07:46 UTC


Leaf-02#show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
```

```
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 00:20:45/00:02:16, RP 10.1.255.4, flags: SJC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan101, Forward/Sparse, 00:20:45/00:02:16, flags:

(*, 239.1.1.2), 00:20:45/00:02:16, RP 10.1.255.4, flags: SJC
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan102, Forward/Sparse, 00:20:45/00:02:16, flags:

(*, 224.0.1.40), 00:21:54/00:02:05, RP 10.1.255.4, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback255, Forward/Sparse, 00:21:53/00:02:05, flags:
```

To return, click Verifying Optimized Layer 2 Multicast with Default MDT in the Underlay .

Outputs to Verify Configuration on VTEP 3

```
Leaf-03#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
172.16.16.1       GigabitEthernet1/0/1  00:12:30/00:01:34 v2    1 / S P G
172.16.26.2       GigabitEthernet1/0/2  00:12:29/00:01:33 v2    1 / S P G



Leaf-03#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-03#show ip mroute 225.0.0.101
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
```

```
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 00:12:43/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:12:36/00:02:23, flags:

(172.16.254.3, 225.0.0.101), 00:12:27/00:02:46, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:12:27/00:02:32, flags:


Leaf-03#show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4

Leaf-03#show ip igmp snooping querier vlan 101
IP address               : 172.16.254.6
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2


Leaf-03#show ip igmp snooping groups vlan 101
Vlan      Group                   Type        Version     Port List
------------------------------------------------------------------------------
101       239.1.1.1               igmp        v2          Gi1/0/10, Tu0




Leaf-03#show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.6:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
```

```
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    MCAST IP:        225.0.0.101
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 0 MAC/IP, 0 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 1 MAC/IP, 0 IMET, 0 EAD


Leaf-03#show l2vpn evpn multicast local address 239.1.1.1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        IGMPv2    N/A     (*, 239.1.1.1)

Leaf-03#show l2vpn evpn multicast remote address 239.1.1.1EVI   VLAN  Originator
 Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
101   101   172.16.254.4        IGMPv2    N/A     (*, 239.1.1.1)



Leaf-03#show l2route evpn multicast routes group 239.1.1.1
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
101   0          239.1.1.1       *               Gi1/0/10:101, V:10101 225.0.0.101


Leaf-03#show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin              Group           Filter     Source(s)
----- ---------- ------------------- --------------- ---------- --------------------
101   0          Gi1/0/10:101        239.1.1.1       N/A        (*)IGMPv2
101   0          172.16.254.4        239.1.1.1       N/A        (*)IGMPv2


Leaf-03#show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 11
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:06:53 UTC
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
```

```
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 5 2022 20:06:52 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 23
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:07:46 UTC


Leaf-03#show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 14
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 5 2022 20:07:19 UTC
```

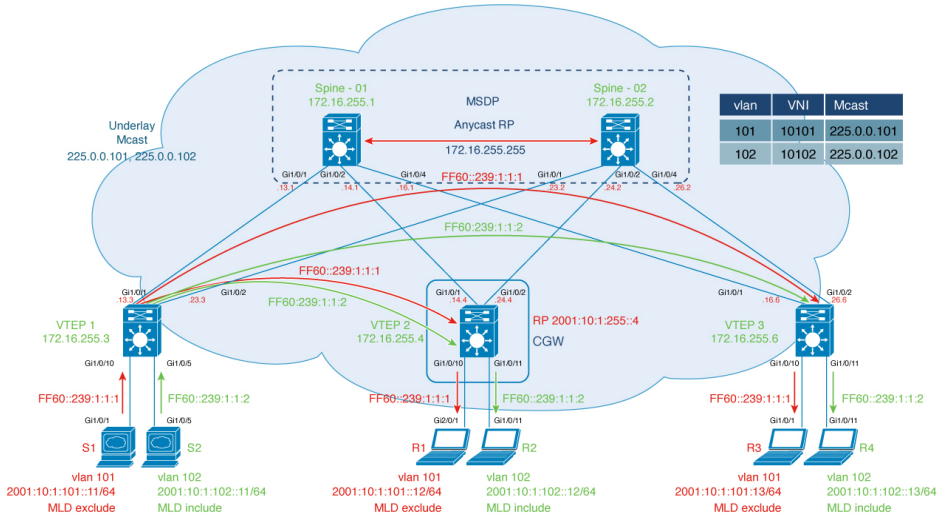
To return, click Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay .

Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication

This example shows how to configure Optimized Layer 2 Multicast for IPv4 and IPv6 multicast traffic across the Layer 2 overlays in an EVPN VXLAN fabric that has underlay multicast replication.

For IPv4 multicast traffic, refer Topology for Optimized Layer 2 Overlay Multicast with Underlay Multicast Replication.

Figure 29: Topology for Optimized Layer 2 Overlay Multicast for IPv6 with Underlay Multicast Replication



The topology shows an EVPN VXLAN network with the source connected to Layer 2 VTEP 1 and receivers connected to Layer 2 VTEP 3. A Centralized Gateway is configured on VTEP 2. The multicast group for VLAN 101 is 225.0.0.101 and for VLAN 102 is 225.0.0.102. The following tables provide sample configurations for the devices in this topology:

Table 33: Configure VTEP 1, CGW, and VTEP 3 for Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 Traffic, with Underlay Multicast Replication

VTEP 1	CGW	VTEP 3

VTEP 1	CGW	VTEP 3
Leaf-01#**show running-config** hostname Leaf-01 ! ip routing ip multicast-routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.3 ip igmp snooping querier ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ! l2vpn evpn replication-type static router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport	Leaf-02#**show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ip multicast-routing ! ipv6 unicast-routing ipv6 mld snooping ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode	Leaf-03#**show running-config** hostname Leaf-03 ! ip routing ip multicast-routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.6 ip igmp snooping querier ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ! l2vpn evpn replication-type static router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6

VTEP 1	CGW	VTEP 3
ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 switchport access vlan 102 switchport mode access ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-address 172.16.255.255 ! end	ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.1.255.4 255.255.255.255 ip pim sparse-mode ipv6 address 2001:10:1:255::4/128 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address 2001:10:1:101::1/64 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ipv6 address 2001:10:1:102::1/64 ! interface Vlan901 no ip address ipv6 enable ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 mcast-group	255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! ip pim rp-addr 172.16.255.255 ! end

VTEP 1	CGW	VTEP 3
	225.0.0.102 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute static redistribute connected exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! ip pim vrf green rp-address 10.1.255.4 ipv6 pim vrf green rp-address 2001:10:1:255::4 ! end	

Table 34: Configure Spine Switch 1 and Spine Switch 2 for Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication

Spine Switch 1	Spine Switch 2

Spine Switch 1	Spine Switch 2
Spine-01#**show running-config**	Spine-02#**show running-config**
hostname Spine-01	hostname Spine-02
!	!
ip routing	ip routing
!	!
ip multicast-routing	ip multicast-routing
!	!
interface Loopback0	interface Loopback0
ip address 172.16.255.1 255.255.255.255	ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface Loopback1	interface Loopback1
ip address 172.16.254.1 255.255.255.255	ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface Loopback2	interface Loopback2
ip address 172.16.255.255 255.255.255.255	ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode	ip pim sparse-mode
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1
no switchport	no switchport
ip address 172.16.13.1 255.255.255.0	ip address 172.16.23.2 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2
no switchport	no switchport
ip address 172.16.14.1 255.255.255.0	ip address 172.16.24.2 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
interface GigabitEthernet1/0/4	interface GigabitEthernet1/0/4
no switchport	no switchport
ip address 172.16.16.1 255.255.255.0	ip address 172.16.26.2 255.255.255.0
ip pim sparse-mode	ip pim sparse-mode
ip ospf network point-to-point	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf 1 area 0
!	!
router ospf 1	router ospf 1
router-id 172.16.255.1	router-id 172.16.255.2
!	!
router bgp 65001	router bgp 65001
bgp router-id 172.16.255.1	bgp router-id 172.16.255.2
bgp log-neighbor-changes	bgp log-neighbor-changes
no bgp default ipv4-unicast	no bgp default ipv4-unicast
neighbor 172.16.255.2 remote-as 65001	neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0	neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.3 remote-as 65001	neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0	neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001	neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0	neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001	neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0	neighbor 172.16.255.6 update-source Loopback0
!	!
address-family ipv4	address-family ipv4
exit-address-family	exit-address-family
!	!
address-family l2vpn evpn	address-family l2vpn evpn
neighbor 172.16.255.2 activate	neighbor 172.16.255.1 activate

Spine Switch 1	Spine Switch 2
neighbor 172.16.255.2 send-community both neighbor 172.16.255.2 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end	neighbor 172.16.255.1 send-community both neighbor 172.16.255.1 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end

To return, click Configuration Example for Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication.

To return to the Configuration Examples section, click Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214

Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic

The following sections provide sample output of the **show** commands to verify Optimized Layer 2 Overlay Multicast with default MDT for IPv4 and IPv6 multicast traffic on the devices in the topology configured above.

To see the output of the **show** commands for IPv4 multicast traffic, go to Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay , on page 264.

The outputs of the **show** commands for IPv6 multicast traffic are provided in the following: sections.

Outputs to verify configuration on VTEP 1

Outputs to verify configuration on CGW

Outputs to verify configuration on VTEP 3

Outputs to verify configuration on Spine Switch 1

Outputs to verify configuration on Spine Switch 2

Outputs to Verify Configuration on VTEP 1

Verify the configuration with MLD exclude

```
Leaf-01# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
172.16.13.1       GigabitEthernet1/0/1    1d00h/00:01:19    v2    1 / S P G
```

```
172.16.23.2       GigabitEthernet1/0/2     1d00h/00:01:34    v2    1 / S P G


Leaf-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-01# show ip mroute 225.0.0.101
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 1d00h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:01:44, flags:

(172.16.254.6, 225.0.0.101), 1d00h/00:01:27, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:17, flags:

(172.16.254.4, 225.0.0.101), 1d00h/00:01:11, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:47, flags:

(172.16.254.3, 225.0.0.101), 1d00h/00:02:58, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:11, flags:


Leaf-01# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 172.16.254.4


Leaf-01#show ipv6 mld snooping querier vlan 101
IP address               : FE80:0:2A56:0:12B3:D5FF:FE6A:8F80
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
```

```
Robustness variable      : 2


Leaf-01# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
----------------------------------------
Total number of channels: 1
Total number of hosts   : 2

Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Tu0       FE80::AC10:FE04                 101  1d00h
1d00h    /
                                                                                       -


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE06                 101  1d00h
1d00h    /
                                                                                       -



Leaf-01# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    MCAST IP:        225.0.0.101
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 3 MAC/IP, 0 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD

Leaf-01# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
```

```
101   101   172.16.254.4         MLDv2     EXCLUDE (*, FF06::239:1:1:1)
101   101   172.16.254.6         MLDv2     EXCLUDE (*, FF06::239:1:1:1)


Leaf-01#show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin              Group                                  Filter
Source(s)
----- ---------- ------------------- -------------------------------------- ----------
--------------------
101   0          172.16.254.4        FF06::239:1:1:1                        EXCLUDE
(*)MLDv2
101   0          172.16.254.6        FF06::239:1:1:1                        EXCLUDE
(*)MLDv2


Leaf-01# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 42
Paths: (1 available, best #1, table evi_101)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 25
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC


Leaf-01# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 20
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
```

```
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:41 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 13
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:41 UTC
```

Verify the configuration with MLD include

```
Leaf-01# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
172.16.13.1       GigabitEthernet1/0/1    1d00h/00:01:34    v2    1 / S P G
172.16.23.2       GigabitEthernet1/0/2    1d00h/00:01:20    v2    1 / S P G


Leaf-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-01# show ip mroute 225.0.0.102
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
```

```
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.102), 1d00h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:08, flags:

(172.16.254.6, 225.0.0.102), 1d00h/00:01:51, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:54, flags:

(172.16.254.4, 225.0.0.102), 1d00h/00:01:43, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:00:09, flags:

(172.16.254.3, 225.0.0.102), 1d00h/00:01:56, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:06, flags:


Leaf-01# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- ------------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d 172.16.254.4
  Y   2001:10:1:102::1                        102   102   7c21.0dbd.954d 172.16.254.4


Leaf-01# show ipv6 mld snooping querier vlan 102
IP address               : FE80:0:2A56:0:12B3:D5FF:FE6A:8F80
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-01# show ipv6 mld snooping membership vlan 102 source 2001:10:1:102::11 group
FF06::239:1:1:2
Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------------


2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE04                 102  1d00h
1d00h    /
                                                                                        -


2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE06                 102  1d00h
1d00h    /
                                                                                        -


Leaf-01# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
```

```
  RD:                172.16.254.3:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    MCAST IP:        225.0.0.102
    Pseudoports:
      GigabitEthernet1/0/5 service instance 102
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 0 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-01# show l2vpn evpn multicast remote address FF06::239:1:1:2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.4         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
102   102   172.16.254.6         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
```

```
Leaf-01# show l2route evpn multicast smet group FF06::239:1:1:2
EVI   ETAG       Origin               Group                                   Filter
Source(s)
----- ---------- -------------------- --------------------------------------- ----------
--------------------
102   0          172.16.254.4         FF06::239:1:1:2                         INCLUDE
2001:10:1:102::11
102   0          172.16.254.6         FF06::239:1:1:2                         INCLUDE
2001:10:1:102::11
```

```
Leaf-01# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 44
Paths: (1 available, best #1, table evi_102)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
```

```
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 31
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC


Leaf-01#show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 22
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:41 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 15
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
```

```
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:41 UTC
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 282.

Outputs to Verify Configuration on CGW

Verify the configuration with MLD exclude.

```
Leaf-02#show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.14.1       GigabitEthernet1/0/1     1d00h/00:01:39    v2    1 / S P G
172.16.24.2       GigabitEthernet1/0/2     1d00h/00:01:30    v2    1 / S P G


Leaf-02#show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-02# show ip mroute 225.0.0.101
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 1d00h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel2, Forward/Sparse-Dense, 1d00h/00:00:12, flags:

(172.16.254.6, 225.0.0.101), 1d00h/00:01:08, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
```

```
  Outgoing interface list:
    Tunnel2, Forward/Sparse-Dense, 1d00h/00:00:39, flags:

(172.16.254.3, 225.0.0.101), 1d00h/00:01:02, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel2, Forward/Sparse-Dense, 1d00h/00:01:12, flags:

(172.16.254.4, 225.0.0.101), 1d00h/00:02:57, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:15, flags:
```

```
Leaf-02# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- ------------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 Vl101
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 Vl101
```

```
Leaf-02# show ipv6 mld snooping querier vlan 101
IP address               : FE80::46D3:CAFF:FE28:6CC1
MLD  version             : v2
Port                     : Gi1/0/10
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2
```

```
Leaf-02# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
----------------------------------------
Total number of channels: 1
Total number of hosts   : 2

Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Gi1/0/10  FE80::46D3:CAFF:FE28:6CC1       101  00:00:00
00:00:01 /
                                                                                00:00:01


::/FF06::239:1:1:1              Tu2       FE80::AC10:FE06                 101  1d00h
1d00h    /
                                                                                -
```

```
Leaf-02# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
```

```
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    MCAST IP:        225.0.0.101
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Enabled (Asymmetric)
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-02# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-02# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator          Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
101   101   172.16.254.6        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-02# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                   Source
       Next-hop(s)
----- ---------- ----------------------------------------
-------------------------------------- --------------------
101   0          FF06::239:1:1:1                         ::
       Gi1/0/10:101, V:10101 225.0.0.101
```

```
Leaf-02# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin              Group                                   Filter
Source(s)
----- ---------- ------------------- --------------------------------------- ----------
--------------------
101   0          Gi1/0/10:101        FF06::239:1:1:1                         EXCLUDE
(*)MLDv2
101   0          172.16.254.6        FF06::239:1:1:1                         EXCLUDE
(*)MLDv2
```

```
Leaf-02# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 26
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
```

```
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:37:21 UTC

Leaf-02# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 44
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 30
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:41 UTC

Leaf-02# show ipv6 mroute vrf green FF06::239:1:1:1
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State
```

```
(*, FF06::239:1:1:1), 1d00h/never, RP 2001:10:1:255::4, flags: SCJ
  Incoming interface: Tunnel5
  RPF nbr: 2001:10:1:255::4
  Immediate Outgoing interface list:
    Vlan101, Forward, 1d00h/never
```

Verify the configuration with MLD include

```
Leaf-02# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
172.16.14.1       GigabitEthernet1/0/1    1d00h/00:01:26    v2    1 / S P G
172.16.24.2       GigabitEthernet1/0/2    1d00h/00:01:17    v2    1 / S P G


Leaf-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-02# show ip mroute 225.0.0.102
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.102), 1d00h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel2, Forward/Sparse-Dense, 1d00h/00:01:02, flags:

(172.16.254.6, 225.0.0.102), 1d00h/00:01:38, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel2, Forward/Sparse-Dense, 1d00h/00:01:43, flags:

(172.16.254.3, 225.0.0.102), 1d00h/00:01:01, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel2, Forward/Sparse-Dense, 1d00h/00:02:18, flags:
```

```
(172.16.254.4, 225.0.0.102), 1d00h/00:02:07, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:06, flags:


Leaf-02# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d Vl102
  Y   2001:10:1:102::1                        102   102   7c21.0dbd.954d Vl102


Leaf-02# show ipv6 mld snooping querier vlan 102
IP address               : FE80::46D3:CAFF:FE28:6CC2
MLD  version             : v2
Port                     : Gi1/0/11
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-02# show ipv6 mld snooping membership vlan 102 source 2001:10:1:102::11 group
FF06::239:1:1:2
Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------------

2001:10:1:102::11/FF06::239:1:1:2 \
                                Gi1/0/11  FE80::46D3:CAFF:FE28:6CC2       102  1d00h
00:01:32 /
                                                                                        -


2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu2       FE80::AC10:FE06                 102  1d00h
1d00h    /
                                                                                        -


Leaf-02# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.4:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan102
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
```

```
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    MCAST IP:        225.0.0.102
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Enabled (Asymmetric)
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-02# show l2vpn evpn multicast local address FF06::239:1:1:2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11        MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
```

```
Leaf-02# show l2vpn evpn multicast remote address FF06::239:1:1:2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.6         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
```

```
Leaf-02# show l2route evpn multicast routes group FF06::239:1:1:2
EVI   ETAG       Group                                   Source
       Next-hop(s)
----- ---------- ----------------------------------------
---------------------------------------- --------------------
102   0          FF06::239:1:1:2                         ::
       Gi1/0/11:102, V:10102 225.0.0.102
```

```
Leaf-02# show l2route evpn multicast smet group FF06::239:1:1:2
EVI   ETAG       Origin               Group                                   Filter
Source(s)
----- ---------- -------------------- --------------------------------------- ----------
--------------------
102   0          Gi1/0/11:102         FF06::239:1:1:2                         INCLUDE
2001:10:1:102::11
102   0          172.16.254.6         FF06::239:1:1:2                         INCLUDE
2001:10:1:102::11
```

```
Leaf-02# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 28
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
```

```
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:37:21 UTC


Leaf-02# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 46
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 32
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:41 UTC

Leaf-02# show ipv6 mroute vrf green FF06::239:1:1:2
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(2001:10:1:102::11, FF06::239:1:1:2), 1d00h/now, flags: SFTI
  Incoming interface: Vlan102
  RPF nbr: 2001:10:1:102::11
  Outgoing interface list: Null
```

To return, click Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 282.

Outputs to verify configuration on VTEP 3

MLD exclude

```
Leaf-03# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.16.1       GigabitEthernet1/0/1 1d00h/00:01:29    v2    1 / S P G
172.16.26.2       GigabitEthernet1/0/2  1d00h/00:01:42    v2    1 / S P G


Leaf-03# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-03# show ip mroute 225.0.0.101
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 1d00h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:43, flags:

(172.16.254.4, 225.0.0.101), 1d00h/00:01:45, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:56, flags:

(172.16.254.6, 225.0.0.101), 1d00h/00:02:59, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:20, flags:

(172.16.254.3, 225.0.0.101), 1d00h/00:01:34, flags: JTx
```

```
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:43, flags:


Leaf-03# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 172.16.254.4


Leaf-03# show ipv6 mld snooping querier vlan 101
IP address               : FE80:0:356:0:E75:BDFF:FE67:EF00
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-03# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
----------------------------------------
Total number of channels: 1
Total number of hosts   : 2

Source/Group                    Interface Reporter                       Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Gi1/0/10  FE80::EEE1:A9FF:FE37:92C1      101  00:00:00
00:00:05 /
                                                                                 00:00:05


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE04                101  1d00h
1d00h    /
                                                                                       -



Leaf-03# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.6:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
```

```
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    MCAST IP:        225.0.0.101
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 3 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-03# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-03# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator          Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
101   101   172.16.254.4        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-03# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                   Source
       Next-hop(s)
----- ---------- --------------------------------------
--------------------------------------- --------------------
101   0          FF06::239:1:1:1                         ::
       Gi1/0/10:101, V:10101 225.0.0.101
```

```
Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin              Group                                   Filter
Source(s)
----- ---------- ------------------- --------------------------------------- ----------
--------------------
101   0          Gi1/0/10:101        FF06::239:1:1:1                         EXCLUDE
(*)MLDv2
101   0          172.16.254.4        FF06::239:1:1:1                         EXCLUDE
(*)MLDv2
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 25
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
```

```
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 42
Paths: (1 available, best #1, table evi_101)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC


Leaf-03# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 6
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:37:34 UTC
```

MLD include

```
Leaf-03# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface               Uptime/Expires    Ver   DR
Address                                                           Prio/Mode
172.16.16.1       GigabitEthernet1/0/1  1d00h/00:01:19    v2    1 / S P G
172.16.26.2       GigabitEthernet1/0/2  1d00h/00:01:32    v2    1 / S P G


Leaf-03# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Leaf-03# show ip mroute 225.0.0.102
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
```

```
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.102), 1d00h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:37, flags:

(172.16.254.6, 225.0.0.102), 1d00h/00:03:28, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:51, flags: A

(172.16.254.3, 225.0.0.102), 1d00h/00:01:08, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:51, flags:

(172.16.254.4, 225.0.0.102), 1d00h/00:01:07, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 1d00h/00:02:50, flags:
```

```
Leaf-03# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                              102   102   7c21.0dbd.954d 172.16.254.4
  Y   2001:10:1:102::1                        102   102   7c21.0dbd.954d 172.16.254.4
```

```
Leaf-03# show ipv6 mld snooping querier vlan 102IP address               :
FE80:0:356:0:E75:BDFF:FE67:EF00
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2
```

```
Leaf-03# show ipv6 mld snooping membership vlan 102 source 2001:10:1:102::11 group
FF06::239:1:1:2
Source/Group                   Interface Reporter                       Vlan Uptime
Last-Join/
                                                                              Last-Leave
-------------------------------------------------------------------------------------------

2001:10:1:102::11/FF06::239:1:1:2 \
                               Gi1/0/11  FE80::EEE1:A9FF:FE37:92C2       102  1d00h
00:00:55 /
```

```
2001:10:1:102::11/FF06::239:1:1:2 \
                                Tu0       FE80::AC10:FE04                 102  1d00h
1d00h    /
```

```
Leaf-03# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.6:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Static (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    MCAST IP:        225.0.0.102
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 0 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 0 IMET, 0 EAD
```

```
Leaf-03# show l2vpn evpn multicast local address FF06::239:1:1:2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11        MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
```

```
Leaf-03# show l2vpn evpn multicast remote address FF06::239:1:1:2
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
102   102   172.16.254.4         MLDv2     INCLUDE (2001:10:1:102::11, FF06::239:1:1:2)
```

```
Leaf-03# show l2route evpn multicast routes group FF06::239:1:1:2
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
---------------------------------------- --------------------
102   0          FF06::239:1:1:2                          ::
       Gi1/0/11:102, V:10102 225.0.0.102
```

```
Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:2
EVI   ETAG       Origin               Group                                   Filter
Source(s)
----- ---------- -------------------- --------------------------------------- ----------
--------------------
102   0          Gi1/0/11:102         FF06::239:1:1:2                         INCLUDE
2001:10:1:102::11
102   0          172.16.254.4         FF06::239:1:1:2                         INCLUDE
2001:10:1:102::11


Leaf-03# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 31
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Apr 7 2022 14:39:04 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51,
 version 44
Paths: (1 available, best #1, table evi_102)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.4]/51
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:38:47 UTC


Leaf-03# show bgp l2vpn evpn route-type 6 0 2001:10:1:102::11 FF06::239:1:1:2 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:102][0][128][2001:10:1:102::11][128][FF06::239:1:1:2][32][172.16.254.6]/51,
 version 8
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
```

```
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on Apr 7 2022 14:37:38 UTC
```

Return to Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 282

Outputs to verify configuration on Spine Switch 1

MLD exclude

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 38, main routing table version 38
37 network entries using 14208 bytes of memory
74 path entries using 16576 bytes of memory
5/5 BGP path/bestpath attribute entries using 1480 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
5 BGP extended community entries using 184 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32568 total bytes of memory
BGP activity 37/0 prefixes, 74/0 paths, scan interval 60 secs
37 networks peaked at 14:48:25 Apr 7 2022 UTC (1d00h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001    1651    1650       38    0    0 1d00h           37
172.16.255.3    4        65001    1639    1658       38    0    0 1d00h            8
172.16.255.4    4        65001    1644    1651       38    0    0 1d00h           19
172.16.255.6    4        65001    1643    1652       38    0    0 1d00h           10


Spine-01# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.13.3       GigabitEthernet1/0/1     1d00h/00:01:36    v2    1 / DR S P G
172.16.14.4       GigabitEthernet1/0/2     1d00h/00:01:25    v2    1 / DR S P G
172.16.16.6       GigabitEthernet1/0/4     1d00h/00:01:41    v2    1 / DR S P G


Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Spine-01# show ip msdp peer
MSDP Peer 172.16.254.2 (?), AS 65001 (configured AS)
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback1 (172.16.254.1)
```

```
    Uptime(Downtime): 1d00h, Messages sent/received: 1599/1600
    Output messages discarded: 0
    Connection and counters cleared 1d00h    ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 4
  Number of connection transitions to Established state: 1
    Input queue size: 0, Output queue size: 0
  MD5 signature protection on MSDP TCP connection: not enabled
  Message counters:
    RPF Failure count: 0
    SA Messages in/out: 1599/1588
    SA Requests in: 0
    SA Responses out: 0
    Data Packets in/out: 4/3


Spine-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 1d00h/00:02:26, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1d00h/00:02:26, flags:

(*, 225.0.0.102), 1d00h/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.102), 1d00h/00:02:37, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(*, 225.0.0.101), 1d00h/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.101), 1d00h/00:01:21, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null
```

MLD include

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 38, main routing table version 38
37 network entries using 14208 bytes of memory
74 path entries using 16576 bytes of memory
5/5 BGP path/bestpath attribute entries using 1480 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
5 BGP extended community entries using 184 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32568 total bytes of memory
BGP activity 37/0 prefixes, 74/0 paths, scan interval 60 secs
37 networks peaked at 14:48:25 Apr 7 2022 UTC (1d00h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001    1661    1660       38    0    0 1d00h           37
172.16.255.3    4        65001    1649    1668       38    0    0 1d00h            8
172.16.255.4    4        65001    1654    1661       38    0    0 1d00h           19
172.16.255.6    4        65001    1653    1662       38    0    0 1d00h           10


Spine-01# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.13.3       GigabitEthernet1/0/1     1d00h/00:01:36    v2    1 / DR S P G
172.16.14.4       GigabitEthernet1/0/2     1d00h/00:01:26    v2    1 / DR S P G
172.16.16.6       GigabitEthernet1/0/4     1d00h/00:01:40    v2    1 / DR S P G


Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)

Spine-01# show ip msdp peer
MSDP Peer 172.16.254.2 (?), AS 65001 (configured AS)
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback1 (172.16.254.1)
    Uptime(Downtime): 1d00h, Messages sent/received: 1610/1610
    Output messages discarded: 0
    Connection and counters cleared 1d00h    ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 4
  Number of connection transitions to Established state: 1
    Input queue size: 0, Output queue size: 0
  MD5 signature protection on MSDP TCP connection: not enabled
  Message counters:
    RPF Failure count: 0
    SA Messages in/out: 1609/1599
```

```
    SA Requests in: 0
    SA Responses out: 0
    Data Packets in/out: 4/3


Spine-01#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 1d00h/00:02:06, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1d00h/00:02:06, flags:

(*, 225.0.0.102), 1d00h/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.102), 1d00h/00:01:34, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(*, 225.0.0.101), 1d00h/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.101), 1d00h/00:01:59, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null
```

Return to Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 282

Outputs to verify configuration on Spine Switch 2

MLD exclude

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 38, main routing table version 38
37 network entries using 14208 bytes of memory
74 path entries using 16576 bytes of memory
5/5 BGP path/bestpath attribute entries using 1480 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
```

```
5 BGP extended community entries using 184 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32568 total bytes of memory
BGP activity 37/0 prefixes, 74/0 paths, scan interval 60 secs
37 networks peaked at 14:48:25 Apr 7 2022 UTC (1d00h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1651    1653       38    0    0 1d00h           37
172.16.255.3    4        65001    1645    1655       38    0    0 1d00h            8
172.16.255.4    4        65001    1648    1656       38    0    0 1d00h           19
172.16.255.6    4        65001    1646    1653       38    0    0 1d00h           10
```

```
Spine-02# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.23.3       GigabitEthernet1/0/1     1d00h/00:01:29    v2    1 / DR S P G
172.16.24.4       GigabitEthernet1/0/2     1d00h/00:01:31    v2    1 / DR S P G
172.16.26.6       GigabitEthernet1/0/4     1d00h/00:01:34    v2    1 / DR S P G
```

```
Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
```

```
Spine-02# show ip msdp peer
MSDP Peer 172.16.254.1 (?), AS 65001 (configured AS)
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback1 (172.16.254.2)
    Uptime(Downtime): 1d00h, Messages sent/received: 1602/1601
    Output messages discarded: 0
    Connection and counters cleared 1d00h    ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 2
  Number of connection transitions to Established state: 1
    Input queue size: 0, Output queue size: 0
  MD5 signature protection on MSDP TCP connection: not enabled
  Message counters:
    RPF Failure count: 0
    SA Messages in/out: 1593/1595
    SA Requests in: 0
    SA Responses out: 0
    Data Packets in/out: 3/6
```

```
Spine-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
```

```
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 1d00h/00:03:18, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:02:45, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:57, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:18, flags:
    Loopback2, Forward/Sparse, 1d00h/00:02:12, flags:

(*, 225.0.0.102), 1d00h/00:03:14, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:02:58, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:09, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:14, flags:

(172.16.254.6, 225.0.0.102), 1d00h/00:01:10, flags: T
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:14, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:09, flags:

(172.16.254.4, 225.0.0.102), 1d00h/00:02:58, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:00, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:14, flags:

(172.16.254.3, 225.0.0.102), 1d00h/00:03:08, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:17, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:13, flags:

(*, 225.0.0.101), 1d00h/00:03:26, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:26, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:20, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:12, flags:

(172.16.254.6, 225.0.0.101), 1d00h/00:03:02, flags: MT
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:12, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:20, flags:

(172.16.254.4, 225.0.0.101), 1d00h/00:02:19, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.24.4
  Outgoing interface list:
```

```
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:26, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:12, flags:

(172.16.254.3, 225.0.0.101), 1d00h/00:02:20, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:26, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:20, flags:
```

MLD include

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 38, main routing table version 38
37 network entries using 14208 bytes of memory
74 path entries using 16576 bytes of memory
5/5 BGP path/bestpath attribute entries using 1480 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
5 BGP extended community entries using 184 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32568 total bytes of memory
BGP activity 37/0 prefixes, 74/0 paths, scan interval 60 secs
37 networks peaked at 14:48:25 Apr 7 2022 UTC (1d00h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1661    1662       38    0    0 1d00h           37
172.16.255.3    4        65001    1654    1664       38    0    0 1d00h            8
172.16.255.4    4        65001    1658    1665       38    0    0 1d00h           19
172.16.255.6    4        65001    1655    1663       38    0    0 1d00h           10


Spine-02# show ip pim neighbor
PIM Neighbor Table
Mode: B - Bidir Capable, DR - Designated Router, N - Default DR Priority,
      P - Proxy Capable, S - State Refresh Capable, G - GenID Capable,
      L - DR Load-balancing Capable
Neighbor          Interface                Uptime/Expires    Ver   DR
Address                                                            Prio/Mode
172.16.23.3       GigabitEthernet1/0/1     1d00h/00:01:16    v2    1 / DR S P G
172.16.24.4       GigabitEthernet1/0/2     1d00h/00:01:17    v2    1 / DR S P G
172.16.26.6       GigabitEthernet1/0/4     1d00h/00:01:21    v2    1 / DR S P G

Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)


Spine-02# show ip msdp peer
MSDP Peer 172.16.254.1 (?), AS 65001 (configured AS)
  Connection status:
    State: Up, Resets: 0, Connection source: Loopback1 (172.16.254.2)
    Uptime(Downtime): 1d00h, Messages sent/received: 1611/1610
    Output messages discarded: 0
    Connection and counters cleared 1d00h    ago
  SA Filtering:
    Input (S,G) filter: none, route-map: none
    Input RP filter: none, route-map: none
    Output (S,G) filter: none, route-map: none
    Output RP filter: none, route-map: none
  SA-Requests:
```

```
    Input filter: none
  Peer ttl threshold: 0
  SAs learned from this peer: 2
  Number of connection transitions to Established state: 1
    Input queue size: 0, Output queue size: 0
  MD5 signature protection on MSDP TCP connection: not enabled
  Message counters:
    RPF Failure count: 0
    SA Messages in/out: 1602/1604
    SA Requests in: 0
    SA Responses out: 0
    Data Packets in/out: 3/6


Spine-02#show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 1d00h/00:03:19, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:02, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:19, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:02:41, flags:
    Loopback2, Forward/Sparse, 1d00h/00:02:38, flags:

(*, 225.0.0.102), 1d00h/00:03:19, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:19, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:30, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:02:37, flags:

(172.16.254.6, 225.0.0.102), 1d00h/00:01:58, flags: T
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:02:59, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:57, flags:

(172.16.254.4, 225.0.0.102), 1d00h/00:03:09, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:23, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:02:53, flags:

(172.16.254.3, 225.0.0.102), 1d00h/00:03:06, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
```

```
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:03:19, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:36, flags:

(*, 225.0.0.101), 1d00h/00:02:46, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:02:46, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:42, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:02:33, flags:

(172.16.254.6, 225.0.0.101), 1d00h/00:01:52, flags: MT
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:02:58, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:08, flags:

(172.16.254.4, 225.0.0.101), 1d00h/00:02:39, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:02:53, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 1d00h/00:03:01, flags:

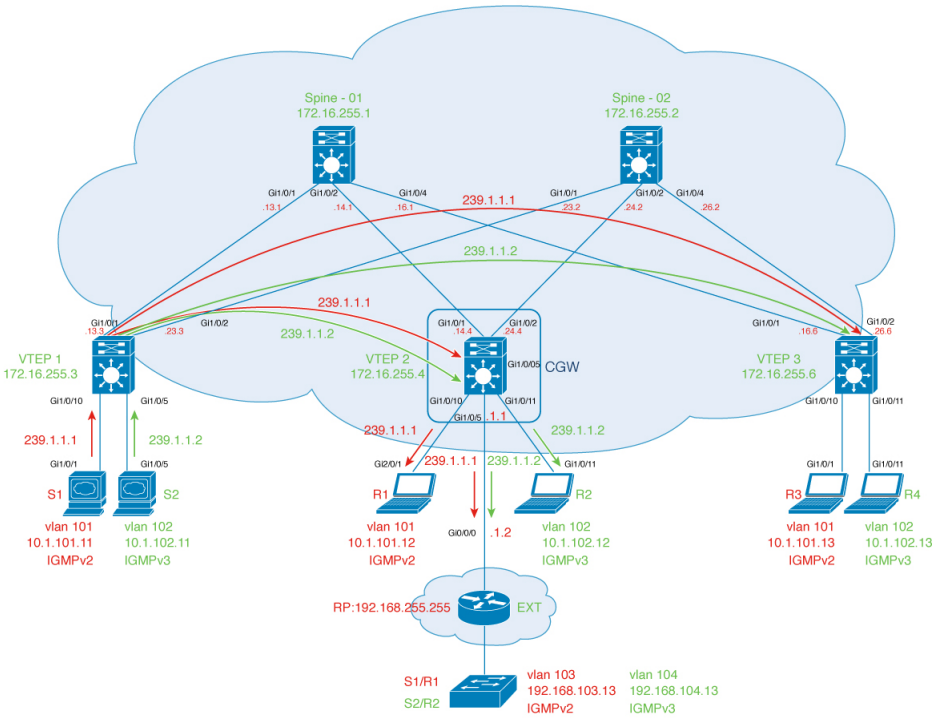
(172.16.254.3, 225.0.0.101), 1d00h/00:02:37, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d00h/00:02:46, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:59, flags:
```

Return to Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 282

Example 5: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4

This example shows how to configure Optimized Layer 2 Overlay Multicast for IPv4 multicast traffic in an EVPN VXLAN fabric that has ingress replication enabled in the underlay and multicast traffic is handedoff to an external network.

Figure 30: Topology Showing Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 Traffic



The topology shows an EVPN VXLAN fabric, with two spine switches and three VTEPs, connected to an external network. Centralized gateway is configured on VTEP 2. The source is connected to VTEP1 and receivers are spread across VTEP 3, centralized gateway, and the external network. VLAN 101 has IGMPv2 and VLAN 102 has IGMPv3 enabled. The rendezvous point (RP) is located outside the EVPN VXLAN fabric. The following tables provide sample configurations for the devices in this topology.

Table 35: Configure VTEP1, CGW, and VTEP3

VTEP1	CGW	VTEP3

VTEP1	CGW	VTEP3
Leaf-01#**show running-config** hostname Leaf-01 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.3 ip igmp snooping querier ! l2vpn evpn replication-type ingress router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 switchport access vlan 102 switchport mode access ! interface GigabitEthernet1/0/10 switchport access vlan 101	Leaf-02#**show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing vrf green ! l2vpn evpn replication-type ingress router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4	Leaf-03#**show running-config** hostname Leaf-03 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.6 ip igmp snooping querier ! l2vpn evpn replication-type ingress router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11

VTEP1	CGW	VTEP3
switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end	255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 no switchport mtu 1500 vrf forwarding green ip address 192.168.1.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 2 vrf green redistribute connected redistribute bgp 65001 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2	switchport access vlan 102 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end

VTEP1	CGW	VTEP3
	remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute static redistribute connected redistribute ospf 2 match internal external 1 external 2 exit-address-family ! ip pim vrf green rp-address 192.168.255.255 ! end	

Table 36: Configure Spine Switch 1, Spine Switch 2, and External Device

Spine Switch 1	Spine Switch 2	External Device
		EXT#**show running-config** hostname EXT ! ip multicast-routing distributed ! bridge-domain 103 bridge-domain 104 ! interface Loopback255 ip address 192.168.255.255 255.255.255.255 ip pim sparse-mode ip ospf 2 area 0 ! interface GigabitEthernet0/0/0 ip address 192.168.1.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 negotiation auto ! interface GigabitEthernet0/0/3 no ip address negotiation auto service instance 103 ethernet encapsulation default bridge-domain 103 ! service instance 104 ethernet encapsulation dot1q 104 rewrite ingress tag pop 1 symmetric bridge-domain 104 ! interface BDI103 ip address 192.168.103.1 255.255.255.0 ip pim sparse-mode ip ospf 2 area 0 ! interface BDI104 ip address 192.168.104.1 255.255.255.0 ip pim sparse-mode ip ospf 2 area 0 ! router ospf 2 passive-interface BDI103 passive-interface BDI104 ! no ip http server no ip http secure-server no ip forward-protocol nd ! ip pim rp-address 192.168.255.255 ! end

Spine Switch 1	Spine Switch 2	External Device
Spine-01#**show running-config**	Spine-02#**show running-config**	
hostname Spine-01	hostname Spine-02	
!	!	
ip routing	ip routing	
!	!	
interface Loopback0	interface Loopback0	
ip address 172.16.255.1	ip address 172.16.255.2	
255.255.255.255	255.255.255.255	
ip ospf 1 area 0	ip ospf 1 area 0	
!	!	
interface Loopback1	interface Loopback1	
ip address 172.16.254.1	ip address 172.16.254.2	
255.255.255.255	255.255.255.255	
ip ospf 1 area 0	ip ospf 1 area 0	
!	!	
interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1	
no switchport	no switchport	
ip address 172.16.13.1	ip address 172.16.23.2	
255.255.255.0	255.255.255.0	
ip ospf network point-to-point	ip ospf network point-to-point	
ip ospf 1 area 0	ip ospf 1 area 0	
!	!	
interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2	
no switchport	no switchport	
ip address 172.16.14.1	ip address 172.16.24.2	
255.255.255.0	255.255.255.0	
ip ospf network point-to-point	ip ospf network point-to-point	
ip ospf 1 area 0	ip ospf 1 area 0	
!	!	
interface GigabitEthernet1/0/4	interface GigabitEthernet1/0/4	
no switchport	no switchport	
ip address 172.16.16.1	ip address 172.16.26.2	
255.255.255.0	255.255.255.0	
ip ospf network point-to-point	ip ospf network point-to-point	
ip ospf 1 area 0	ip ospf 1 area 0	
!	!	
router ospf 1	router ospf 1	
router-id 172.16.255.1	router-id 172.16.255.2	
!	!	
router bgp 65001	router bgp 65001	
bgp router-id 172.16.255.1	bgp router-id 172.16.255.2	
bgp log-neighbor-changes	bgp log-neighbor-changes	
no bgp default ipv4-unicast	no bgp default ipv4-unicast	
neighbor 172.16.255.2	neighbor 172.16.255.1	
remote-as 65001	remote-as 65001	
neighbor 172.16.255.2	neighbor 172.16.255.1	
update-source Loopback0	update-source Loopback0	
neighbor 172.16.255.3	neighbor 172.16.255.3	
remote-as 65001	remote-as 65001	
neighbor 172.16.255.3	neighbor 172.16.255.3	
update-source Loopback0	update-source Loopback0	
neighbor 172.16.255.4	neighbor 172.16.255.4	
remote-as 65001	remote-as 65001	
neighbor 172.16.255.4	neighbor 172.16.255.4	
update-source Loopback0	update-source Loopback0	
neighbor 172.16.255.6	neighbor 172.16.255.6	
remote-as 65001	remote-as 65001	
neighbor 172.16.255.6	neighbor 172.16.255.6	
update-source Loopback0	update-source Loopback0	
!	!	
address-family ipv4	address-family ipv4	

Spine Switch 1	Spine Switch 2	External Device
exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.2 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! end	exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.1 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! end	

To return, click Example 5: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4, on page 313.

To return to the Configuration Examples section, click Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214.

Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4

The following sections provide sample output of **show** commands to verify Optimized Layer 2 Overlay Multicast configuration for IPv4 multicast traffic handoff to an external network, on the devices in the topology configured above.

Outputs to verify configuration on VTEP 1

Outputs to verify configuration on Centralized Gateway

Outputs to verify configuration on VTEP 3

Outputs to verify configuration on External Device

Outputs to verify configuration on Spine Switch 1

Outputs to verify configuration on Spine Switch 2

Outputs to Verify Configuration on VTEP 1

Verify IGMPv2

```
Leaf-01# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                EVI   VLAN  MAC Address    Source
```

```
----- ------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                            101   101   7c21.0dbd.9541 172.16.254.4
```

```
Leaf-01# show ip igmp snooping querier vlan 101
IP address               : 172.16.254.3
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2
```

```
Leaf-01# show ip igmp snooping groups vlan 101
Vlan      Group                Type        Version     Port List
-------------------------------------------------------------------------
101       239.1.1.1            igmp        v2          Tu0
```

```
Leaf-01# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-01# show l2vpn evpn multicast local address 239.1.1.1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
```

```
Leaf-01# show l2vpn evpn multicast remote address 239.1.1.1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         IGMPv2    N/A     (*, 239.1.1.1)
101   101   172.16.254.6         IGMPv2    N/A     (*, 239.1.1.1)
```

```
Leaf-01# show l2route evpn multicast routes group 239.1.1.1
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
101   0          239.1.1.1       *               V:10101 172.16.254.4, V:10101 172.16.254.6


Leaf-01# show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin                Group           Filter     Source(s)
----- ---------- --------------------- --------------- ---------- --------------------
101   0          172.16.254.4          239.1.1.1       N/A        (*)IGMPv2
101   0          172.16.254.6          239.1.1.1       N/A        (*)IGMPv2

Leaf-01# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 33
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from [6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 10:23:03 UTC
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 17
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 10:23:15 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 10:22:59 UTC

Leaf-01# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 245
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from [6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
```

```
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 14:12:06 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 243
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 14:12:06 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 14:12:06 UTC
```

Verify IGMPv3

```
Leaf-01# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                  EVI   VLAN  MAC Address    Source
----- ---------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                               102   102   7c21.0dbd.954d 172.16.254.4

Leaf-01#show ip igmp snooping querier vlan 102
IP address               : 172.16.254.3
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2


Leaf-01# show ip igmp snooping groups vlan 102
Vlan      Group                Type        Version     Port List
-------------------------------------------------------------------------
102       239.1.1.2            igmp        v2,v3       Tu0


Leaf-01# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.3:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
```

```
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    Pseudoports:
      GigabitEthernet1/0/5 service instance 102
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD


Leaf-01# show l2vpn evpn multicast local address 239.1.1.2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------


Leaf-01# show l2vpn evpn multicast remote address 239.1.1.2
EVI   VLAN  Originator          Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
102   102   172.16.254.4        IGMPv2    N/A     (*, 239.1.1.2)
102   102   172.16.254.4        IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)
102   102   172.16.254.6        IGMPv2    N/A     (*, 239.1.1.2)
102   102   172.16.254.6        IGMPv3    INCLUDE (192.168.104.13, 239.1.1.2)


Leaf-01# show l2route evpn multicast routes group 239.1.1.2
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
102   0          239.1.1.2       *              V:10102 172.16.254.4, V:10102 172.16.254.6


Leaf-01# show l2route evpn multicast smet group 239.1.1.2
EVI   ETAG       Origin              Group           Filter     Source(s)
----- ---------- ------------------- --------------- ---------- --------------------
102   0          172.16.254.4        239.1.1.2       INCLUDE    (*)IGMPv2 10.1.102.11
102   0          172.16.254.6        239.1.1.2       INCLUDE    (*)IGMPv2 192.168.104.13


Leaf-01# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 48
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:24:12 UTC
BGP routing table entry for
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 46
```

```
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 12:24:12 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:24:12 UTC


Leaf-01# show bgp l2vpn evpn route-type 6 0 192.168.104.13 239.1.1.2 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27, version
122
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:25:57 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27, version
121
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 12:25:57 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:2
```

To return, click Verifying the Configuration of Optimized Layer 2 Multicast.

Outputs to Verify Configuration on Centralized Gateway

Verify IGMPv2

```
Leaf-02# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 Vl101


Leaf-02# show ip igmp snooping querier vlan 101
IP address               : 10.1.101.1
IGMP version             : v2
Port                     : Router
Max response time        : 10s


Leaf-02#show ip igmp snooping groups vlan 101
Vlan      Group                   Type        Version     Port List
------------------------------------------------------------------------
101       239.1.1.1               igmp        v2          Gi1/0/10, Tu0


Leaf-02# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD


Leaf-02# show l2vpn evpn multicast local address 239.1.1.1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
```

```
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        IGMPv2    N/A     (*, 239.1.1.1)


Leaf-02# show l2route evpn multicast routes group 239.1.1.1
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
101   0          239.1.1.1       *               Gi1/0/10, V:10101 172.16.254.6



Leaf-02# show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin               Group           Filter     Source(s)
----- ---------- -------------------- --------------- ---------- --------------------
101   0          Gi1/0/10:101         239.1.1.1       N/A        (*)IGMPv2

Leaf-02# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 16
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 10:20:56 UTC


Leaf-02# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 208
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from [6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23
 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 14:12:06 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 206
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
```

```
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 14:12:06 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 14:12:06 UTC


Leaf-02# show ip mroute vrf green 239.1.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 05:51:44/stopped, RP 192.168.255.255, flags: SJCF
  Incoming interface: GigabitEthernet1/0/5, RPF nbr 192.168.1.2
  Outgoing interface list:
    Vlan101, Forward/Sparse, 05:51:44/00:02:14, flags:

(192.168.103.13, 239.1.1.1), 00:11:18/00:01:47, flags: JT
  Incoming interface: GigabitEthernet1/0/5, RPF nbr 192.168.1.2
  Outgoing interface list:
    Vlan101, Forward/Sparse, 00:11:18/00:02:14, flags:

(10.1.101.11, 239.1.1.1), 01:49:45/00:01:58, flags: FT
  Incoming interface: Vlan101, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/5, Forward/Sparse, 00:10:52/00:03:02, flags:

Leaf-02#show ip igmp vrf green groups 239.1.1.1 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source,
       Ac - Group accounted towards access control limit

Interface:      Vlan101
Group:          239.1.1.1
Flags:
Uptime:         05:51:45
Group mode:     EXCLUDE (Expires: 00:02:14)
Last reporter:  10.1.101.12
Source list is empty
```

Verify IPGMv3

```
Leaf-02# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                EVI   VLAN  MAC Address    Source
----- ------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                            102   102   7c21.0dbd.954d Vl102

Leaf-02# show ip igmp snooping querier vlan 102
IP address               : 10.1.102.1
IGMP version             : v3
Port                     : Router
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2


Leaf-02# show ip igmp snooping groups vlan 102
Vlan      Group                   Type        Version     Port List
------------------------------------------------------------------------
102       239.1.1.2               igmp        v2,v3       Gi1/0/11, Tu0


Leaf-02# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.4:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan102
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.4
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD


Leaf-02# show l2vpn evpn multicast local address 239.1.1.2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
```

```
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11        IGMPv2    N/A     (*, 239.1.1.2)
102   102   Gi1/0/11        IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)

Leaf-02# show l2vpn evpn multicast remote address 239.1.1.2
EVI   VLAN  Originator          Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
102   102   172.16.254.6        IGMPv2    N/A     (*, 239.1.1.2)
102   102   172.16.254.6        IGMPv3    INCLUDE (192.168.104.13, 239.1.1.2)


Leaf-02# show l2route evpn multicast routes group 239.1.1.2
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
102   0          239.1.1.2       *               Gi1/0/11:102, V:10102 172.16.254.6



Leaf-02# show l2route evpn multicast smet group 239.1.1.2
EVI   ETAG       Origin              Group           Filter     Source(s)
----- ---------- ------------------- --------------- ---------- --------------------
102   0          Gi1/0/11:102        239.1.1.2       INCLUDE    (*)IGMPv2 10.1.102.11
102   0          172.16.254.6        239.1.1.2       INCLUDE    (*)IGMPv2 192.168.104.13

Leaf-02# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 50
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:24:12 UTC


Leaf-02# show bgp l2vpn evpn route-type 6 0 192.168.104.13 239.1.1.2 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27, version
93
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.6:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:25:57 UTC
BGP routing table entry for
```

```
[6][172.16.254.6:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27, version
91
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 12:25:57 UTC
  Refresh Epoch 1
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:25:57 UTC


Leaf-02# show ip mroute vrf green 239.1.1.2
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(10.1.102.11, 239.1.1.2), 00:14:59/00:03:28, flags: sTI
  Incoming interface: Vlan102, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/5, Forward/Sparse, 00:08:53/00:03:28, flags:

(192.168.104.13, 239.1.1.2), 00:15:06/00:02:53, flags: sTI
  Incoming interface: GigabitEthernet1/0/5, RPF nbr 192.168.1.2
  Outgoing interface list:
    Vlan102, Forward/Sparse, 00:15:06/00:02:53, flags:


Leaf-02# show ip igmp vrf green groups 239.1.1.1 detail

Flags: L - Local, U - User, SG - Static Group, VG - Virtual Group,
       SS - Static Source, VS - Virtual Source,
       Ac - Group accounted towards access control limit

Interface:      Vlan101
Group:          239.1.1.1
Flags:
```

```
Uptime:         06:44:06
Group mode:     EXCLUDE (Expires: 00:02:57)
Last reporter:  10.1.101.12
Source list is empty
Leaf-02#!
```

To return, click Verifying the Configuration of Optimized Layer 2 Multicast .

Outputs to Verify Configuration on VTEP 3

Verify IGMPv2

```
Leaf-03# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                  EVI   VLAN  MAC Address    Source
----- ---------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                               101   101   7c21.0dbd.9541 172.16.254.4

Leaf-03# show ip igmp snooping querier vlan 101
IP address               : 172.16.254.6
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2

Leaf-03# show ip igmp snooping groups vlan 101
Vlan      Group                Type        Version     Port List
-------------------------------------------------------------------------
101       239.1.1.1            igmp        v2          Gi1/0/10, Tu0

Leaf-03# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.6:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 1 MAC/IP
```

```
    Peers:
      172.16.254.3
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD

Leaf-03# show l2vpn evpn multicast local address 239.1.1.1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        IGMPv2    N/A     (*, 239.1.1.1)



Leaf-03# show l2vpn evpn multicast remote address 239.1.1.1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         IGMPv2    N/A     (*, 239.1.1.1)


Leaf-03# show l2route evpn multicast routes group 239.1.1.1
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
101   0          239.1.1.1       *               V:10101 172.16.254.4, Gi1/0/10:101


Leaf-03# show l2route evpn multicast smet group 239.1.1.1
EVI   ETAG       Origin               Group           Filter     Source(s)
----- ---------- -------------------- --------------- ---------- --------------------
101   0          Gi1/0/10:101         239.1.1.1       N/A        (*)IGMPv2
101   0          172.16.254.4         239.1.1.1       N/A        (*)IGMPv2

Leaf-03# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 10
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 10:23:15 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 10:22:59 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23, version 25
Paths: (1 available, best #1, table evi_101)
  Flag: 0x100
  Not advertised to any peer
  Refresh Epoch 1
```

```
  Local, imported path from [6][172.16.254.4:101][0][0][*][32][239.1.1.1][32][172.16.254.4]/23
 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 10:22:59 UTC
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 * 239.1.1.1 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][32][239.1.1.1][32][172.16.254.6]/23, version 125
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 14:12:06 UTC
```

Verify IGMPv3

```
Leaf-03# show l2vpn evpn default-gateway vlan 102
Valid Default Gateway Address                  EVI   VLAN  MAC Address    Source
----- ---------------------------------------- ----- ----- -------------- -----------
  Y   10.1.102.1                               102   102   7c21.0dbd.954d 172.16.254.4
```

```
Leaf-03# show ip igmp snooping querier vlan 102
IP address               : 172.16.254.6
IGMP version             : v3
Port                     : Switch
Max response time        : 10s
Query interval           : 60s
Robustness variable      : 2
```

```
Leaf-03# show ip igmp snooping groups vlan 102
Vlan      Group                   Type        Version     Port List
-------------------------------------------------------------------------
102       239.1.1.2               igmp        v2,v3       Gi1/0/11, Tu0
```

```
Leaf-03# show l2vpn evpn evi 102 detail
EVPN instance:       102 (VLAN Based)
  RD:                172.16.254.6:102 (auto)
  Import-RTs:        65001:102
  Export-RTs:        65001:102
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
```

```
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              102
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10102
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    Pseudoports:
      GigabitEthernet1/0/11 service instance 102
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-03# show l2vpn evpn multicast local address 239.1.1.2
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
102   102   Gi1/0/11        IGMPv2    N/A     (*, 239.1.1.2)
102   102   Gi1/0/11        IGMPv3    INCLUDE (192.168.104.13, 239.1.1.2)
```

```
Leaf-03# show l2vpn evpn multicast remote address 239.1.1.2
EVI   VLAN  Originator          Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
102   102   172.16.254.4        IGMPv2    N/A     (*, 239.1.1.2)
102   102   172.16.254.4        IGMPv3    INCLUDE (10.1.102.11, 239.1.1.2)
```

```
Leaf-03# show l2route evpn multicast routes group 239.1.1.2
EVI   ETAG       Group           Source          Next-hop(s)
----- ---------- --------------- --------------- --------------------
102   0          239.1.1.2       *               Gi1/0/11:102, V:10102 172.16.254.4
```

```
Leaf-03# show l2route evpn multicast smet group 239.1.1.2
EVI   ETAG       Origin              Group           Filter     Source(s)
----- ---------- ------------------- --------------- ---------- --------------------
102   0          Gi1/0/11:102        239.1.1.2       INCLUDE    (*)IGMPv2 192.168.104.13
102   0          172.16.254.4        239.1.1.2       INCLUDE    (*)IGMPv2 10.1.102.11
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 10.1.102.11 239.1.1.2 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 48
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
```

```
      rx pathid: 0, tx pathid: 0
      Updated on May 6 2022 12:24:12 UTC
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:24:12 UTC
BGP routing table entry for
[6][172.16.254.6:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27, version 50
Paths: (1 available, best #1, table evi_102)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from
[6][172.16.254.4:102][0][32][10.1.102.11][32][239.1.1.2][32][172.16.254.4]/27 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:24:12 UTC
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 192.168.104.13 239.1.1.2 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:102][0][32][192.168.104.13][32][239.1.1.2][32][172.16.254.6]/27, version
90
Paths: (1 available, best #1, table evi_102)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v3
      Extended Community: RT:65001:102 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on May 6 2022 12:25:57 UTC
```

To return, click Verifying the Configuration of Optimized Layer 2 Multicast .

Outputs to Verify Configuration on the External Device

Verify IGMPv2

```
EXT# show ip igmp groups 239.1.1.1
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter   Group Accounted
239.1.1.1        BDI103                     00:34:05  00:02:59  192.168.103.13
```

```
EXT# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 192.168.255.255 (?)


EXT# show ip mroute 239.1.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 03:22:07/00:02:45, RP 192.168.255.255, flags: SJCF
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    BDI103, Forward/Sparse, 00:34:20/00:02:43, flags:
    GigabitEthernet0/0/0, Forward/Sparse, 03:22:07/00:02:45, flags:

(192.168.103.13, 239.1.1.1), 00:34:46/00:03:24, flags: FT
  Incoming interface: BDI103, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 00:34:20/00:03:16, flags:

(10.1.101.11, 239.1.1.1), 02:13:13/00:02:10, flags: T
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.1.1
  Outgoing interface list:
    BDI103, Forward/Sparse, 00:34:20/00:02:43, flags:
```

Verify IGMPv3

```
EXT# show ip igmp groups 239.1.1.2
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter   Group Accounted
239.1.1.2        BDI104                   00:20:51  stopped   192.168.104.13


EXT# show ip pim rp mappingPIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 192.168.255.255 (?)



EXT# show ip mroute 239.1.1.2
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
```

```
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag, l - LISP decap ref count contributor
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
                          t - LISP transit group
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(192.168.104.13, 239.1.1.2), 00:19:53/00:03:14, flags: sT
  Incoming interface: BDI104, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 00:13:40/00:03:14, flags:

(10.1.102.11, 239.1.1.2), 00:20:34/00:02:20, flags: sTI
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 192.168.1.1
  Outgoing interface list:
    BDI104, Forward/Sparse, 00:13:40/00:02:20, flags:
```

To return, click Verifying the Configuration of Optimized Layer 2 Multicast .

Outputs to Verify Configuration on Spine Switch 1

Verify IGMPv2

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 144, main routing table version 144
34 network entries using 13056 bytes of memory
65 path entries using 14560 bytes of memory
15/14 BGP path/bestpath attribute entries using 4440 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
12 BGP extended community entries using 484 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32660 total bytes of memory
BGP activity 55/21 prefixes, 141/76 paths, scan interval 60 secs
35 networks peaked at 13:38:38 May 6 2022 UTC (02:54:38.599 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001     498     499      144    0    0 06:11:11       31
172.16.255.3    4        65001     426     499      144    0    0 06:11:24        6
172.16.255.4    4        65001     433     498      144    0    0 06:11:23       19
172.16.255.6    4        65001     437     498      144    0    0 06:11:26        9
```

Verify IGMPv3

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 144, main routing table version 144
```

```
34 network entries using 13056 bytes of memory
65 path entries using 14560 bytes of memory
15/14 BGP path/bestpath attribute entries using 4440 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
12 BGP extended community entries using 484 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32660 total bytes of memory
BGP activity 56/22 prefixes, 143/78 paths, scan interval 60 secs
35 networks peaked at 13:38:38 May 6 2022 UTC (03:30:10.672 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001     538     538      144    0    0 06:46:43        31
172.16.255.3    4        65001     464     538      144    0    0 06:46:56         6
172.16.255.4    4        65001     474     537      144    0    0 06:46:55        19
172.16.255.6    4        65001     475     537      144    0    0 06:46:58         9
```

To return, click Verifying the Configuration of Optimized Layer 2 Multicast .

Outputs to Verify Configuration on Spine Switch 2

Verify IGMPv2

```
Spine-02# show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 144, main routing table version 144
34 network entries using 13056 bytes of memory
65 path entries using 14560 bytes of memory
15/14 BGP path/bestpath attribute entries using 4440 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
12 BGP extended community entries using 484 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32660 total bytes of memory
BGP activity 53/19 prefixes, 141/76 paths, scan interval 60 secs
35 networks peaked at 13:38:38 May 6 2022 UTC (02:55:51.914 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     501     499      144    0    0 06:12:24        31
172.16.255.3    4        65001     429     500      144    0    0 06:12:20         6
172.16.255.4    4        65001     434     500      144    0    0 06:12:20        19
172.16.255.6    4        65001     442     500      144    0    0 06:12:23         9
```

Verify IGMPv3

```
Spine-02#show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 144, main routing table version 144
34 network entries using 13056 bytes of memory
65 path entries using 14560 bytes of memory
15/14 BGP path/bestpath attribute entries using 4440 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
12 BGP extended community entries using 484 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 32660 total bytes of memory
BGP activity 55/21 prefixes, 143/78 paths, scan interval 60 secs
35 networks peaked at 13:38:38 May 6 2022 UTC (03:30:25.802 ago)
```

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     539     538      144    0    0 06:46:58        31
172.16.255.3    4        65001     467     538      144    0    0 06:46:54         6
172.16.255.4    4        65001     474     537      144    0    0 06:46:54        19
172.16.255.6    4        65001     480     538      144    0    0 06:46:57         9
```

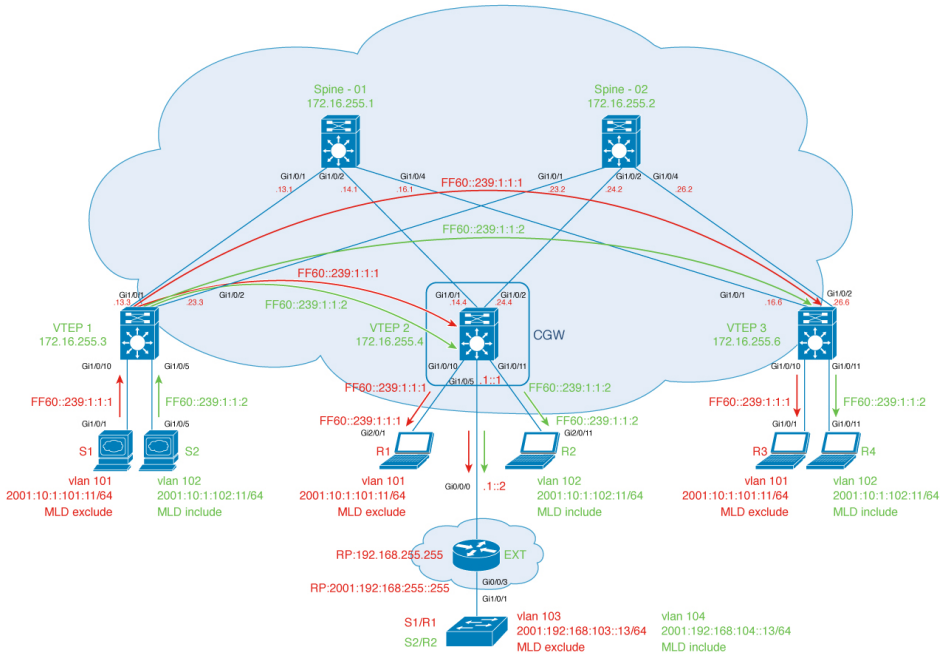
To return, click Verifying the Configuration of Optimized Layer 2 Multicast .

Example 6: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6

This example shows how to configure Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 traffic in an EVPN VXLAN fabric that has ingress replication enabled in the underlay and multicast traffic is handedoff to an external network.

For IPv4 multicast traffic, refer Topology for Optimized Layer 2 Overlay Multicast Interworking with External Network.

Figure 31: Topology for Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv6 Traffic



The topology shows an EVPN VXLAN fabric, with two spine switches and three VTEPs, connected to an external network. Multicast traffic source is located in Layer 2 VTEP 1 and receivers are located in Layer 2 VTEP 3, the Centralized Gateway, and also located in the external network. Centralized gateway (CGW) is configured on VTEP 2. VLAN 101 has IGMPv2 and MLD enabled and VLAN 102 has IGMPv3 and MLD enabled. Rendezvous Point (RP) is located outside the EVPN VXLAN fabric. The following tables provide sample configurations for the devices in this topology:

Table 37: Configure VTEP1, CGW, and VTEP3

VTEP1	CGW	VTEP3

VTEP1	CGW	VTEP3
Leaf-01#**show running-config** hostname Leaf-01 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.3 ip igmp snooping querier ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ! l2vpn evpn replication-type ingress router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5	Leaf-02#**show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ipv6 multicast-routing vrf green ! l2vpn evpn replication-type ingress router-id Loopback1 default-gateway advertise multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 !	Leaf-03#**show running-config** hostname Leaf-03 ! ip routing ! ip igmp snooping querier version 3 ip igmp snooping querier address 172.16.254.6 ip igmp snooping querier ! ipv6 unicast-routing ipv6 mld snooping querier version 2 ipv6 mld snooping querier ipv6 mld snooping ! l2vpn evpn replication-type ingress router-id Loopback1 multicast advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface

VTEP1	CGW	VTEP3
switchport access vlan 102 switchport mode access ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end	interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 no switchport mtu 1500 vrf forwarding green ip address 192.168.1.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ipv6 address 2001:192:168:1::1/64 ospfv3 2 network point-to-point ospfv3 2 ipv6 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address 2001:10:1:101::1/64 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ipv6 address 2001:10:1:102::1/64 !	GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport mode access ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! end

VTEP1	CGW	VTEP3
	interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 10102 ingress-replication ! router ospfv3 2 ! address-family ipv6 unicast vrf green redistribute connected redistribute bgp 65001 exit-address-family ! router ospf 2 vrf green redistribute connected redistribute bgp 65001 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute static redistribute connected redistribute ospf 2 match internal external 1 external 2 exit-address-family ! ip pim vrf green rp-address 192.168.255.255 ! ipv6 pim vrf green rp-address 2001:192:168:255::255	

VTEP1	CGW	VTEP3
	! end	

Table 38: Configure Spine Switch 1, Spine Switch 2, and External Device

Spine Switch 1	Spine Switch 2	External Device

Spine Switch 1	Spine Switch 2	External Device
Spine-01#**show running-config**	Spine-02#**show running-config**	EXT1#**show running-config**
hostname Spine-01	hostname Spine-02	hostname EXT
!	!	!
ip routing	ip routing	ip multicast-routing
!	!	distributed
interface Loopback0	interface Loopback0	!
ip address 172.16.255.1	ip address 172.16.255.2	ipv6 unicast-routing
255.255.255.255	255.255.255.255	ipv6 multicast-routing
ip ospf 1 area 0	ip ospf 1 area 0	!
!	!	bridge-domain 103
interface Loopback1	interface Loopback1	bridge-domain 104
ip address 172.16.254.1	ip address 172.16.254.2	!
255.255.255.255	255.255.255.255	interface Loopback255
ip ospf 1 area 0	ip ospf 1 area 0	ip address 192.168.255.255
!	!	255.255.255.255
interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1	ip pim sparse-mode
no switchport	no switchport	ip ospf 2 area 0
ip address 172.16.13.1	ip address 172.16.23.2	ipv6 address
255.255.255.0	255.255.255.0	2001:192:168:255::255/128
ip ospf network point-to-point	ip ospf network point-to-point	ipv6 ospf 2 area 0
ip ospf 1 area 0	ip ospf 1 area 0	!
!	!	interface GigabitEthernet1/0/0
interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2	ip address 192.168.1.2
no switchport	no switchport	255.255.255.0
ip address 172.16.14.1	ip address 172.16.24.2	ip pim sparse-mode
255.255.255.0	255.255.255.0	ip ospf network point-to-point
ip ospf network point-to-point	ip ospf network point-to-point	ip ospf 2 area 0
ip ospf 1 area 0	ip ospf 1 area 0	negotiation auto
!	!	ipv6 address
interface GigabitEthernet1/0/4	interface GigabitEthernet1/0/4	2001:192:168:1::2/64
no switchport	no switchport	ipv6 ospf 2 area 0
ip address 172.16.16.1	ip address 172.16.26.2	ipv6 ospf network
255.255.255.0	255.255.255.0	point-to-point
ip ospf network point-to-point	ip ospf network point-to-point	!
ip ospf 1 area 0	ip ospf 1 area 0	interface GigabitEthernet1/0/3
!	!	no ip address
router ospf 1	router ospf 1	negotiation auto
router-id 172.16.255.1	router-id 172.16.255.2	service instance 103 ethernet
!	!	encapsulation dot1q 103
router bgp 65001	router bgp 65001	rewrite ingress tag pop 1
bgp router-id 172.16.255.1	bgp router-id 172.16.255.2	symmetric
bgp log-neighbor-changes	bgp log-neighbor-changes	bridge-domain 103
no bgp default ipv4-unicast	no bgp default ipv4-unicast	!
neighbor 172.16.255.2	neighbor 172.16.255.1	service instance 104 ethernet
remote-as 65001	remote-as 65001	encapsulation dot1q 104
neighbor 172.16.255.2	neighbor 172.16.255.1	rewrite ingress tag pop 1
update-source Loopback0	update-source Loopback0	symmetric
neighbor 172.16.255.3	neighbor 172.16.255.3	bridge-domain 104
remote-as 65001	remote-as 65001	!
neighbor 172.16.255.3	neighbor 172.16.255.3	!
update-source Loopback0	update-source Loopback0	interface BDI103
neighbor 172.16.255.4	neighbor 172.16.255.4	ip address 192.168.103.1
remote-as 65001	remote-as 65001	255.255.255.0
neighbor 172.16.255.4	neighbor 172.16.255.4	ip pim sparse-mode
update-source Loopback0	update-source Loopback0	ip ospf 2 area 0
neighbor 172.16.255.6	neighbor 172.16.255.6	ipv6 address
remote-as 65001	remote-as 65001	2001:192:168:103::1/64
neighbor 172.16.255.6	neighbor 172.16.255.6	ipv6 ospf 2 area 0
update-source Loopback0	update-source Loopback0	!
!	!	interface BDI104
address-family ipv4	address-family ipv4	ip address 192.168.104.1
exit-address-family	exit-address-family	255.255.255.0

Spine Switch 1	Spine Switch 2	External Device
!	!	ip pim sparse-mode
address-family l2vpn evpn	address-family l2vpn evpn	ip igmp static-group 239.1.1.2
neighbor 172.16.255.2 activate	neighbor 172.16.255.1 activate	source 10.1.102.11
neighbor 172.16.255.2	neighbor 172.16.255.1	ip igmp version 3
send-community both	send-community both	ip ospf 2 area 0
neighbor 172.16.255.2	neighbor 172.16.255.1	ipv6 address
route-reflector-client	route-reflector-client	2001:192:168:104::1/64
neighbor 172.16.255.3 activate	neighbor 172.16.255.3 activate	ipv6 ospf 2 area 0
neighbor 172.16.255.3	neighbor 172.16.255.3	!
send-community both	send-community both	router ospfv3 2
neighbor 172.16.255.3	neighbor 172.16.255.3	!
route-reflector-client	route-reflector-client	address-family ipv6 unicast
neighbor 172.16.255.4 activate	neighbor 172.16.255.4 activate	passive-interface BDI103
neighbor 172.16.255.4	neighbor 172.16.255.4	passive-interface BDI104
send-community both	send-community both	passive-interface Loopback255
neighbor 172.16.255.4	neighbor 172.16.255.4	exit-address-family
route-reflector-client	route-reflector-client	!
neighbor 172.16.255.6 activate	neighbor 172.16.255.6 activate	router ospf 2
neighbor 172.16.255.6	neighbor 172.16.255.6	passive-interface BDI103
send-community both	send-community both	passive-interface BDI104
neighbor 172.16.255.6	neighbor 172.16.255.6	!
route-reflector-client	route-reflector-client	ip pim rp-address
exit-address-family	exit-address-family	192.168.255.255
!	!	ip pim ssm range SSM
end	end	!
		ip access-list standard SSM
		10 permit 239.1.1.2
		!
		ipv6 pim rp-address
		2001:192:168:255::255
		!
		end

To return, click Example 6: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 341.

To return to the Configuration Examples section, click Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 214.

Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6

The following sections provide sample output of **show** commands to verify Optimized Layer 2 Overlay Multicast configuration for IPv4 and IPv6 multicast traffic handoff to an external network, on the devices in the topology configured above.

To see the output of **show** commands for IPv4 multicast traffic, go to Verify the configuration for IPv4 multicast traffic handoff.

The output of **show** commands for IPv6 multicast traffic are provided in the following sections.

Outputs to verify configuration on VTEP1

Outputs to verify configuration on CGW

Outputs to verify configuration on VTEP3

Outputs to verify configuration on Spine Switch 1

Outputs to verify configuration on Spine Switch 2

Outputs to verify configuration on External Device

Outputs to Verify Configuration on VTEP1

```
Leaf-01# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 172.16.254.4


Leaf-01# show ipv6 mld snooping querier vlan 101
IP address               : FE80:0:7155:0:12B3:D5FF:FE6A:8F80
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2

Leaf-01# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
-----------------------------------------
Total number of channels: 1
Total number of hosts   : 2

Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Tu0       FE80::AC10:FE04                 101  20:58:46
20:58:46 /
                                                                                       -


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE06                 101  20:48:20
20:48:20 /


Leaf-01# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
```

```
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.3
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD

Leaf-01# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------

Leaf-01# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         MLDv2     EXCLUDE (*, FF06::239:1:1:1)
101   101   172.16.254.6         MLDv2     EXCLUDE (*, FF06::239:1:1:1)

Leaf-01# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                   Source
       Next-hop(s)
----- ---------- ----------------------------------------
---------------------------------------- --------------------
101   0          FF06::239:1:1:1                         ::
       V:10101 172.16.254.4, Gi1/0/10:101, V:10101 172.16.254.6

Leaf-01# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin               Group                                   Filter
Source(s)
----- ---------- -------------------- --------------------------------------- ----------
--------------------
101   0          172.16.254.4         FF06::239:1:1:1                         EXCLUDE
(*)MLDv2
101   0          172.16.254.6         FF06::239:1:1:1                         EXCLUDE
(*)MLDv2

Leaf-01# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 655
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:09:15 UTC
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 653
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
```

```
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 9 2022 16:09:15 UTC
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:09:15 UTC

Leaf-01# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.3:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 690
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:19:41 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 688
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 9 2022 16:19:41 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:19:41 UTC
```

To return, click Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 349.

Outputs to Verify Configuration on CGW

```
Leaf-02# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 Vl101
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 Vl101


Leaf-02# show ipv6 mld snooping querier vlan 101
IP address               : FE80::46D3:CAFF:FE28:6CC1
MLD  version             : v2
Port                     : Gi1/0/10
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2


Leaf-02# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
----------------------------------------
Total number of channels: 1
Total number of hosts   : 2

Source/Group                    Interface Reporter                        Vlan Uptime
Last-Join/
                                                                                Last-Leave
-------------------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Gi1/0/10  FE80::46D3:CAFF:FE28:6CC1       101  00:00:00
00:01:28 /
                                                                                 00:01:28


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE06                 101  20:52:37
20:52:37 /
                                                                                 -


Leaf-02# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.4
```

```
    VRF:
    IPv4 IRB:        Enabled (Asymmetric)
    IPv6 IRB:        Enabled (Asymmetric)
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.6
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-02# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-02# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator          Version   Filter  (Source, Group)
----- ----- ------------------- --------- ------- ---------------
101   101   172.16.254.6        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-02# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
--------------------------------------- --------------------
101   0          FF06::239:1:1:1                          ::
       Gi1/0/10:101, V:10101 172.16.254.3, V:10101 172.16.254.6
```

```
Leaf-02# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin              Group                                    Filter
Source(s)
----- ---------- ------------------- ---------------------------------------- ----------
--------------------
101   0          Gi1/0/10:101        FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
101   0          172.16.254.6        FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
```

```
Leaf-02# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 471
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.4
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:09:15 UTC
```

```
Leaf-02# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 513
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35 (global)
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:19:41 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 511
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 9 2022 16:19:41 UTC
  Refresh Epoch 2
  Local
    172.16.254.6 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.6, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:19:41 UTC


Leaf-02# show ipv6 mroute vrf green FF06::239:1:1:1
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06::239:1:1:1), 21:03:04/never, RP 2001:192:168:255::255, flags: SCJ
  Incoming interface: GigabitEthernet1/0/5
  RPF nbr: FE80::A2B4:39FF:FE21:9182
  Immediate Outgoing interface list:
    Vlan101, Forward, 21:03:04/never

(2001:10:1:101::11, FF06::239:1:1:1), 00:05:00/00:02:12, flags: SFJT
  Incoming interface: Vlan101
  RPF nbr: 2001:10:1:101::11
  Immediate Outgoing interface list:
    GigabitEthernet1/0/5, Forward, 00:05:00/00:02:35
```

```
Leaf-02# show ipv6 mld vrf green groups ff06::239:1:1:1 detail
Interface:      Vlan101
Group:          FF06::239:1:1:1
Uptime:         21:03:04
Router mode:    EXCLUDE (Expires: 00:02:49)
Host mode:      INCLUDE
Last reporter:  FE80::7E21:DFF:FEBD:9500
Source list is empty
```

To return, click Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 349.

Outputs to Verify Configuration on VTEP3

```
Leaf-03# show l2vpn evpn default-gateway vlan 101
Valid Default Gateway Address                 EVI   VLAN  MAC Address    Source
----- --------------------------------------- ----- ----- -------------- -----------
  Y   10.1.101.1                              101   101   7c21.0dbd.9541 172.16.254.4
  Y   2001:10:1:101::1                        101   101   7c21.0dbd.9541 172.16.254.4
```

```
Leaf-03# show ipv6 mld snooping querier vlan 101
IP address               : FE80:0:1C56:0:E75:BDFF:FE67:EF00
MLD  version             : v2
Port                     : Switch
Max response time        : 10s
Query interval           : 125s
Robustness variable      : 2
```

```
Leaf-03# show ipv6 mld snooping membership vlan 101
Snooping Membership Summary for Vlan 101
----------------------------------------
Total number of channels: 1
Total number of hosts   : 2

Source/Group                    Interface Reporter                       Vlan Uptime
Last-Join/
                                                                               Last-Leave
-------------------------------------------------------------------------------------------------

::/FF06::239:1:1:1              Gi1/0/10  FE80::EEE1:A9FF:FE37:92C1      101  00:00:00
00:00:10 /
                                                                                 00:00:10


::/FF06::239:1:1:1              Tu0       FE80::AC10:FE04                101  21:05:03
21:05:03 /
                                                                                        -
```

```
Leaf-03# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.6:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
```

```
  Replication Type:  Ingress (global)
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Disabled (global)
  Re-originate RT5:  Disabled
  Adv. Multicast:    Enabled (global)
  Vlan:              101
    Protected:       False
    Ethernet-Tag:    0
    State:           Established
    Flood Suppress:  Attached
    Core If:
    Access If:
    NVE If:          nve1
    RMAC:            0000.0000.0000
    Core Vlan:       0
    L2 VNI:          10101
    L3 VNI:          0
    VTEP IP:         172.16.254.6
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 2 MAC/IP
    Peers:
      172.16.254.3
        Routes: 1 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
        Routes: 2 MAC, 4 MAC/IP, 1 IMET, 0 EAD
```

```
Leaf-03# show l2vpn evpn multicast local address FF06::239:1:1:1
EVI   VLAN  Interface       Version   Filter  (Source, Group)
----- ----- --------------- --------- ------- ---------------
101   101   Gi1/0/10        MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-03# show l2vpn evpn multicast remote address FF06::239:1:1:1
EVI   VLAN  Originator           Version   Filter  (Source, Group)
----- ----- -------------------- --------- ------- ---------------
101   101   172.16.254.4         MLDv2     EXCLUDE (*, FF06::239:1:1:1)
```

```
Leaf-03# show l2route evpn multicast routes group FF06::239:1:1:1
EVI   ETAG       Group                                    Source
       Next-hop(s)
----- ---------- ---------------------------------------
--------------------------------------- --------------------
101   0          FF06::239:1:1:1                          ::
       V:10101 172.16.254.4, V:10101 172.16.254.3, Gi1/0/10:101
```

```
Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:1
EVI   ETAG       Origin               Group                                    Filter
Source(s)
----- ---------- -------------------- ---------------------------------------- ----------
--------------------
101   0          Gi1/0/10:101         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
101   0          172.16.254.4         FF06::239:1:1:1                          EXCLUDE
(*)MLDv2
```

```
Leaf-03# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.4
BGP routing table entry for
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 859
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
```

```
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on May 9 2022 16:09:15 UTC
  Refresh Epoch 2
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:09:15 UTC
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35, version 861
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from
[6][172.16.254.4:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.4]/35 (global)
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:09:15 UTC


Leaf-03# show bgp l2vpn evpn route-type 6 0 * FF06::239:1:1:1 172.16.254.6
BGP routing table entry for
[6][172.16.254.6:101][0][0][*][128][FF06::239:1:1:1][32][172.16.254.6]/35, version 870
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.6)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      IGMP/MLD v2, exclude
      Extended Community: RT:65001:101 ENCAP:8
      Local irb vxlan vtep:
        vrf:not found, l3-vni:0
        local router mac:0000.0000.0000
        core-irb interface:(not found)
        vtep-ip:172.16.254.6
      rx pathid: 0, tx pathid: 0x0
      Updated on May 9 2022 16:19:41 UTC
```

To return, click Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 349.

Outputs to Verify Configuration on Spine Switch 1

```
Spine-01# show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
```

```
BGP table version is 330, main routing table version 330
41 network entries using 15744 bytes of memory
79 path entries using 17696 bytes of memory
13/12 BGP path/bestpath attribute entries using 3848 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
8 BGP extended community entries using 324 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 37732 total bytes of memory
BGP activity 126/85 prefixes, 384/305 paths, scan interval 60 secs
45 networks peaked at 15:44:30 May 9 2022 UTC (21:34:21.817 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001    1698    1699      330    0    0 23:12:41        38
172.16.255.3    4        65001    1567    1710      330    0    0 23:12:44         8
172.16.255.4    4        65001    1624    1710      330    0    0 23:12:44        22
172.16.255.6    4        65001    1548    1718      330    0    0 23:12:44        11
```

To return, click Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 349.

Outputs to Verify Configuration on Spine Switch 2

```
Spine-02# show bgp l2vpn evpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 330, main routing table version 330
41 network entries using 15744 bytes of memory
79 path entries using 17696 bytes of memory
13/12 BGP path/bestpath attribute entries using 3848 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
8 BGP extended community entries using 324 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 37732 total bytes of memory
BGP activity 130/89 prefixes, 386/307 paths, scan interval 60 secs
45 networks peaked at 15:44:30 May 9 2022 UTC (21:34:42.736 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1699    1698      330    0    0 23:13:02        38
172.16.255.3    4        65001    1566    1709      330    0    0 23:13:02         8
172.16.255.4    4        65001    1621    1707      330    0    0 23:13:02        22
172.16.255.6    4        65001    1554    1712      330    0    0 23:13:02        11
```

To return, click Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 349.

Outputs to Verify Configuration on External Device

```
EXT# show ipv6 mld groups  FF06::239:1:1:1
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06::239:1:1:1                         BDI103
   21:07:51  00:03:45
```

```
EXT# show ipv6 pim group-map FF06::239:1:1:1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: 2001:192:168:255::255
    RPF: Tu4,2001:192:168:255::255 (us)
    Info source: Static
    Uptime: 20:57:37, Groups: 1

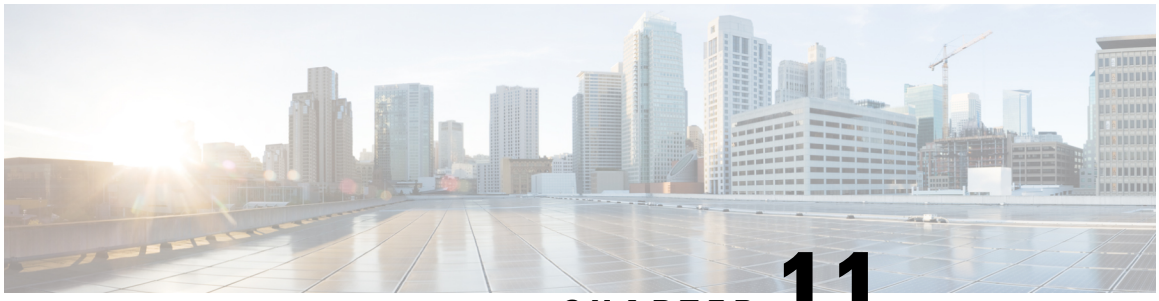

EXT# show ipv6 mroute FF06::239:1:1:1
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06::239:1:1:1), 21:07:51/00:02:36, RP 2001:192:168:255::255, flags: SCJ
  Incoming interface: Tunnel4
  RPF nbr: 2001:192:168:255::255
  Immediate Outgoing interface list:
    BDI103, Forward, 21:07:51/never
    GigabitEthernet0/0/0, Forward, 20:57:13/00:02:36

(2001:10:1:101::11, FF06::239:1:1:1), 00:16:12/00:03:00, RP 2001:192:168:255::255, flags:
SRJ
  Incoming interface: Tunnel4
  RPF nbr: 2001:192:168:255::255
  Immediate Outgoing interface list:
    GigabitEthernet0/0/0, Null, 00:16:12/00:02:36
  Inherited Outgoing interface list:
    BDI103, Forward, 21:07:51/never

(2001:10:1:101::11, FF06::239:1:1:1), 00:16:12/00:03:00, flags: SJT
  Incoming interface: GigabitEthernet0/0/0
  RPF nbr: FE80::7E21:DFF:FEBD:9564
  Inherited Outgoing interface list:
    BDI103, Forward, 21:07:51/never
```

To return, click Verifying Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6, on page 349.



CHAPTER 11

Configuring Tenant Routed Multicast

- Restrictions for Tenant Routed Multicast , on page 361
- Information about Tenant Routed Multicast, on page 361
- How to Configure Tenant Routed Multicast, on page 373
- Troubleshooting Tenant Routed Multicast, on page 566

Restrictions for Tenant Routed Multicast

- Data MDT between BGP EVPN VXLAN TRM and external MVPN multicast network is not supported.
- In the underlay network, the default MDT supports only PIM sparse mode, and Data MDT supports PIM sparse mode and PIM source specific multicast (SSM) mode.

Note The spine switch can be configured as the rendezvous point (RP). The RP for the underlay network can also be elected using PIM Bootstrap Router (BSR) or Auto-RP or Static.

- TRM data MDT is supported for TRM in PIM sparse mode and TRM in PIM source specific mode for IPv4 connections only.
- TRM data MDT supports IPv4 C-multicast flows only.
- TRM data MDT supports rate-based switchover only for IPv4 multicast flows.
- TRM data MDT supports Layer 2, IP, and VRF-Lite handoffs only.

Information about Tenant Routed Multicast

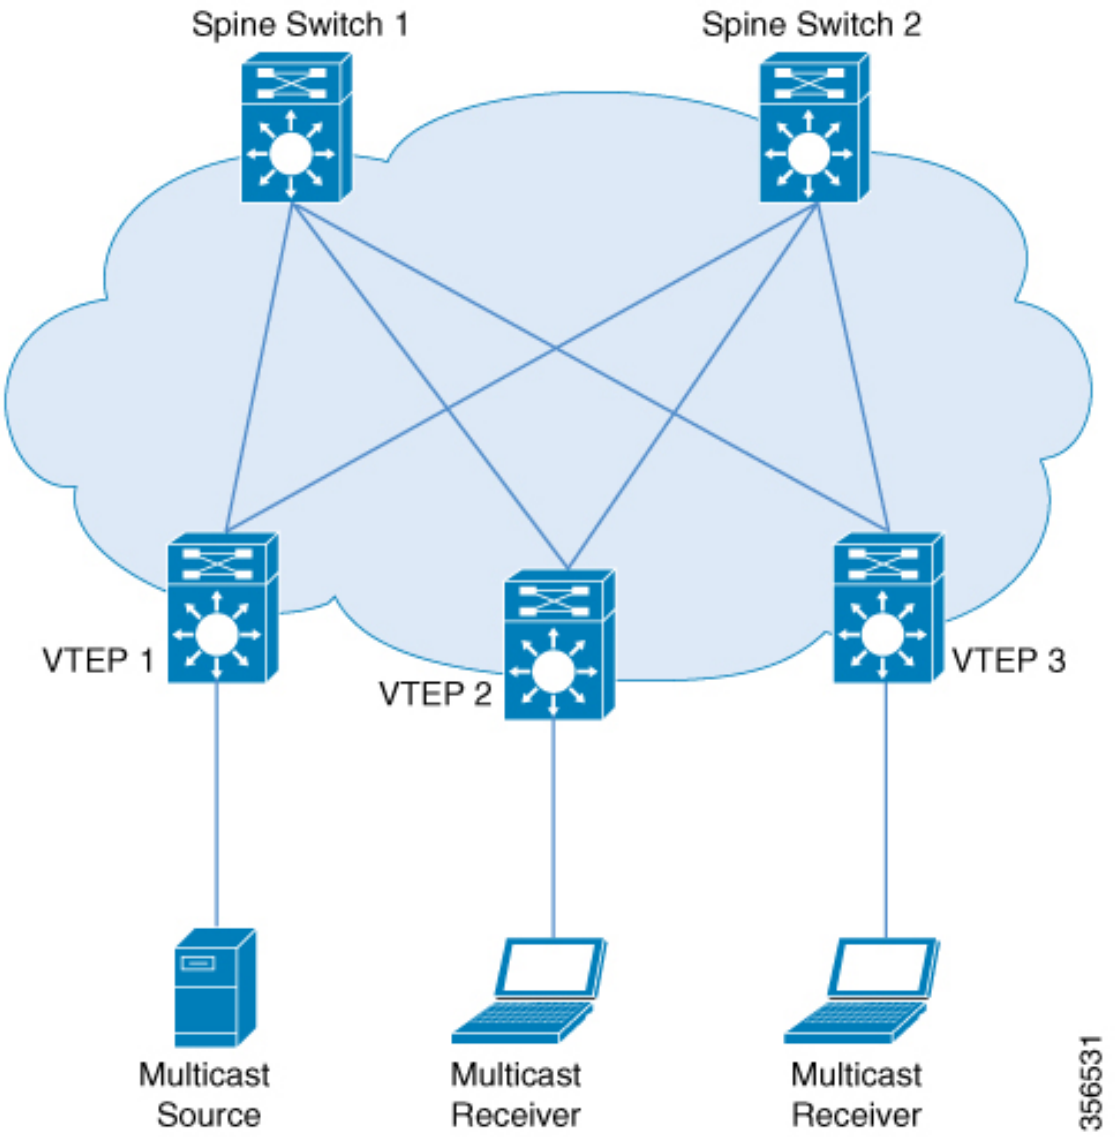
TRM enables multicast forwarding in a VXLAN fabric that uses a BGP-based EVPN control plane. TRM provides multi-tenancy aware multicast forwarding between senders and receivers within the same or different subnets local or across VTEPs.

TRM enables the delivery of a customer's IP multicast traffic in a multi-tenant fabric in an efficient and resilient manner. The delivery of TRM improves Layer 3 overlay multicast functionality in the networks. With TRM enabled, multicast forwarding in the underlay is leveraged to replicate VXLAN-encapsulated routed multicast

traffic. A default MDT is built per-VRF. This is in addition to the existing multicast groups for broadcast and unknown unicast traffic in a Layer 2 virtual network instance (VNI), and for Layer 2 multicast replication group. The individual multicast group addresses in the overlay are mapped to the respective underlay multicast address for replication and transport. The advantage of using a BGP-based approach is that it allows the BGP EVPN VXLAN fabric with TRM to operate as fully distributed overlay rendezvous point (RP), with the RP presence on every edge device or VTEP.

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and multicast rendezvous points, might reside inside the data center but might also be inside the campus or externally reachable via the WAN. Thus, TRM allows a seamless integration with existing multicast networks with newer enterprise fabric.

Figure 32: Tenant Routed Multicast Topology



For IPv4 and IPv6 multicast traffic, TRM uses BGP EVPN and multicast virtual private network (MVPN) routes to perform multicast routing. All the VTEPs in the network do not need to be BGP peers. There can be BGP peering between the VTEPs and the spine switches with the spine switches acting as route reflectors. Source reachability is distributed via EVPN route type 2 and EVPN route type 5 in the fabric. RPF is installed based on these routes. Source-active and receiver-join information is carried in the MVPN address family using route types 5, 6, and 7.

In an EVPN VXLAN network, TRM is supported in the overlay network in PIM sparse mode and PIM source specific multicast (SSM) mode. VTEPs have BGP peering in MVPN as well as EVPN address families to exchange routes for TRM.

TRM in PIM Sparse Mode

PIM-SM distributes information about active sources by forwarding data packets on the shared tree. Because PIM-SM uses shared trees, it requires the use of a rendezvous point (RP). An RP is used for the initial convergence of multicast traffic between sources and receivers.

The following section describes the different ways the RP can be configured for TRM in PIM sparse mode.

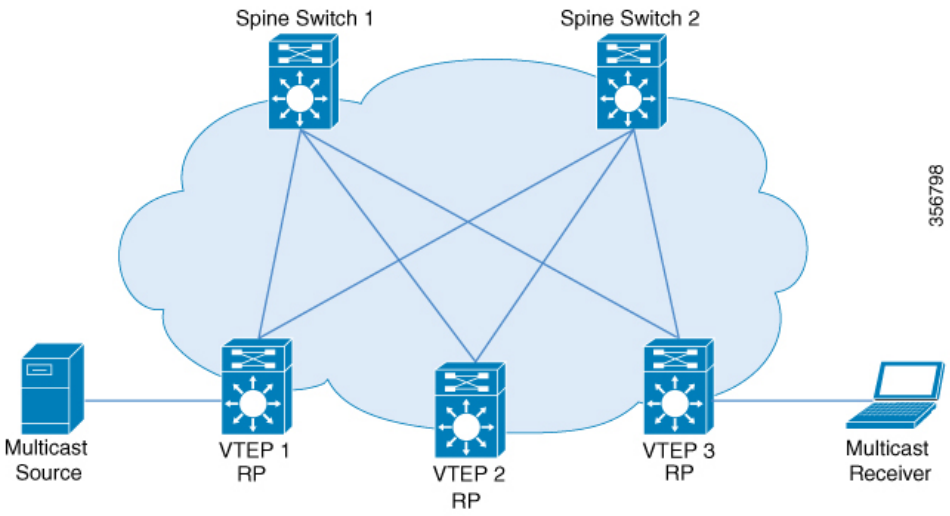
RP Placement

For TRM in PIM sparse mode, the overlay RP can be configured either within the BGP EVPN VXLAN fabric, or outside of the fabric.

Anycast RP in the Overlay Network

In Anycast RP in the overlay network, every VTEP acts as an RP.

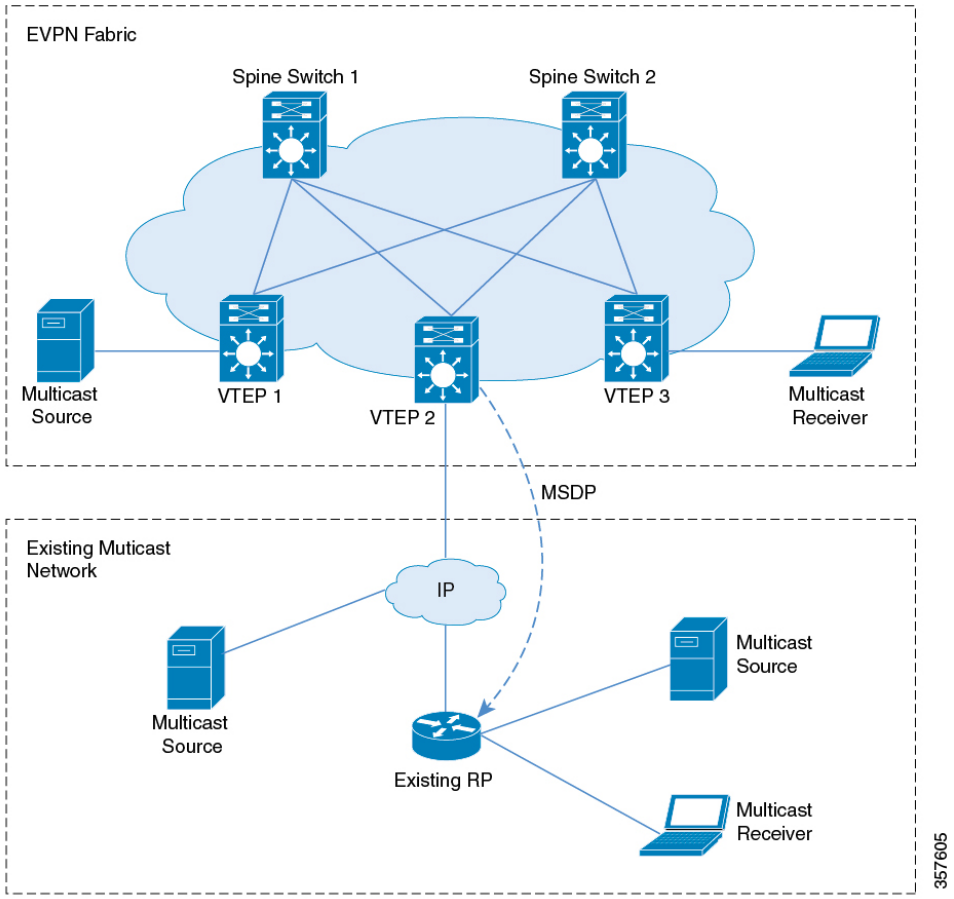
Figure 33: Anycast RP in the Overlay Network



Anycast RP can also inter-operate with existing multicast networks and RP as shown below. Refer the *Configuring MSDP* chapter of the *IP Multicast Routing Configuration Guide* for information on how to configure MSDP.

Note MSDP is an IPv4 protocol and the MSDP originator IP loopback configuration is mandatory.

Figure 34: Anycast RP with IPv4 Overlay in an Existing Multicast Network



RP Inside the BGP EVPN VXLAN Fabric

In a TRM scenario where all sources and receivers are within the EVPN VXLAN network, the overlay RP can be placed on a border spine or on a VTEP.

Figure 35: Border Spine Switch as an RP

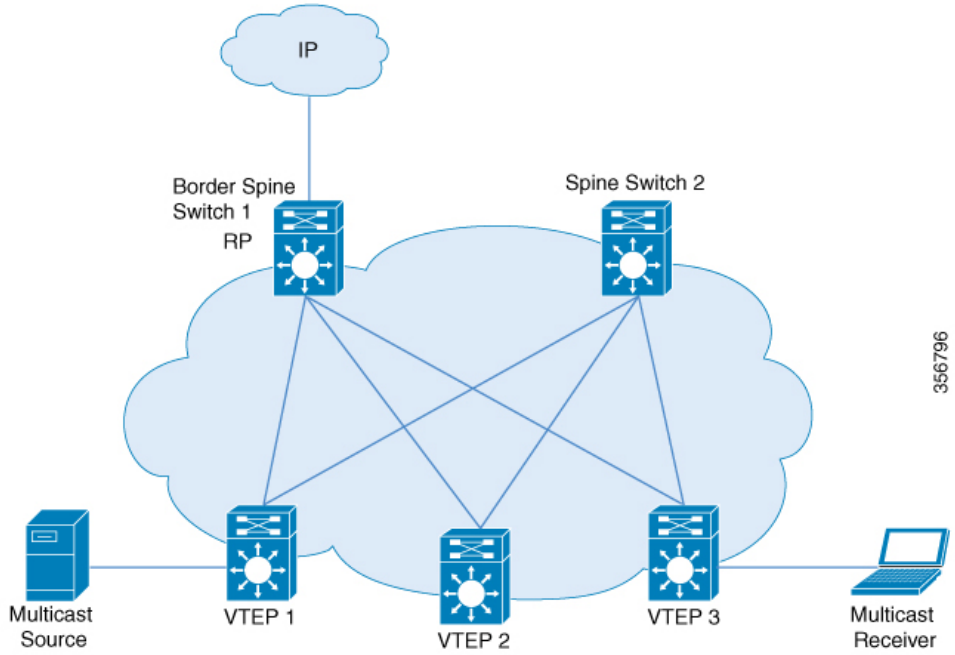
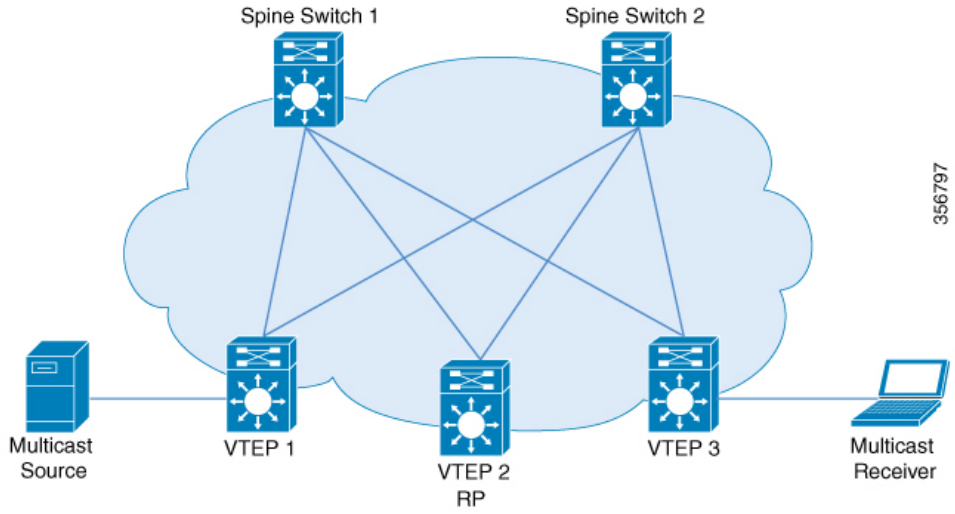


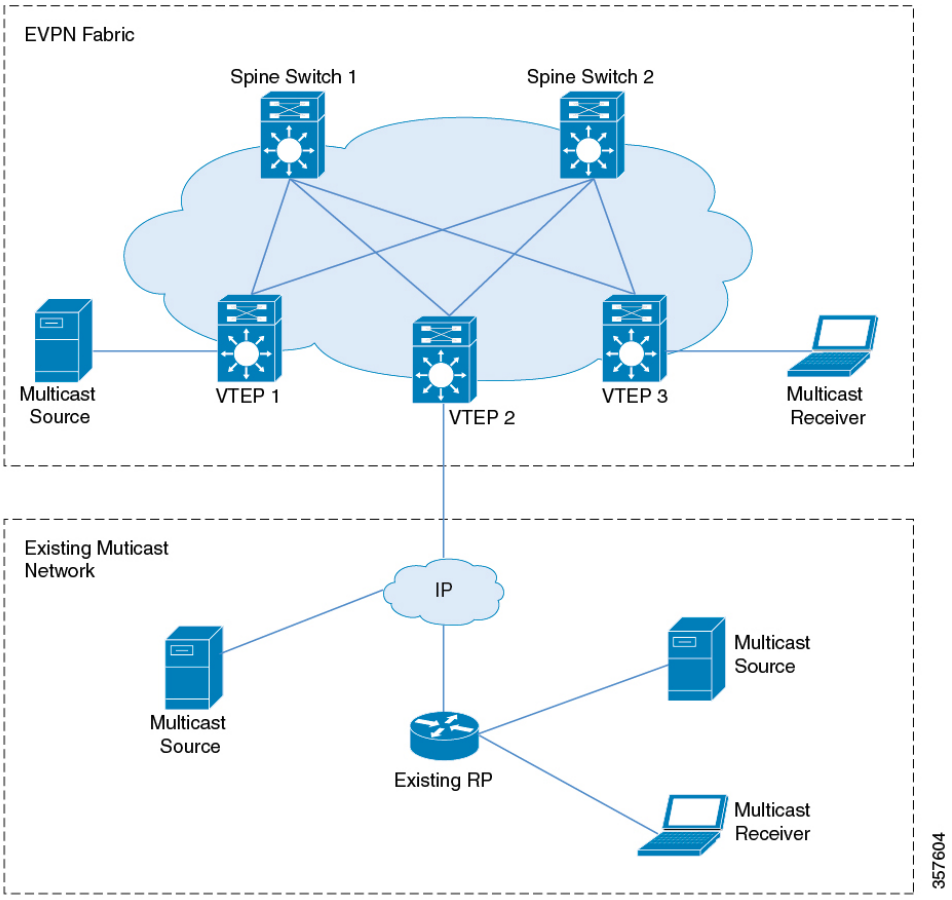
Figure 36: VTEP as an RP



RP Outside the BGP EVPN VXLAN Fabric

In a TRM scenario where sources and receivers need to interwork with the overlay network and the EVPN VXLAN network, the RP can be placed on an external router connected to a VTEP.

Figure 37: RP Outside the BGP EVPN VXLAN Fabric with an Existing Multicast Network



In PIM sparse mode, TRM can be configured in three different ways depending on how the RP is configured:

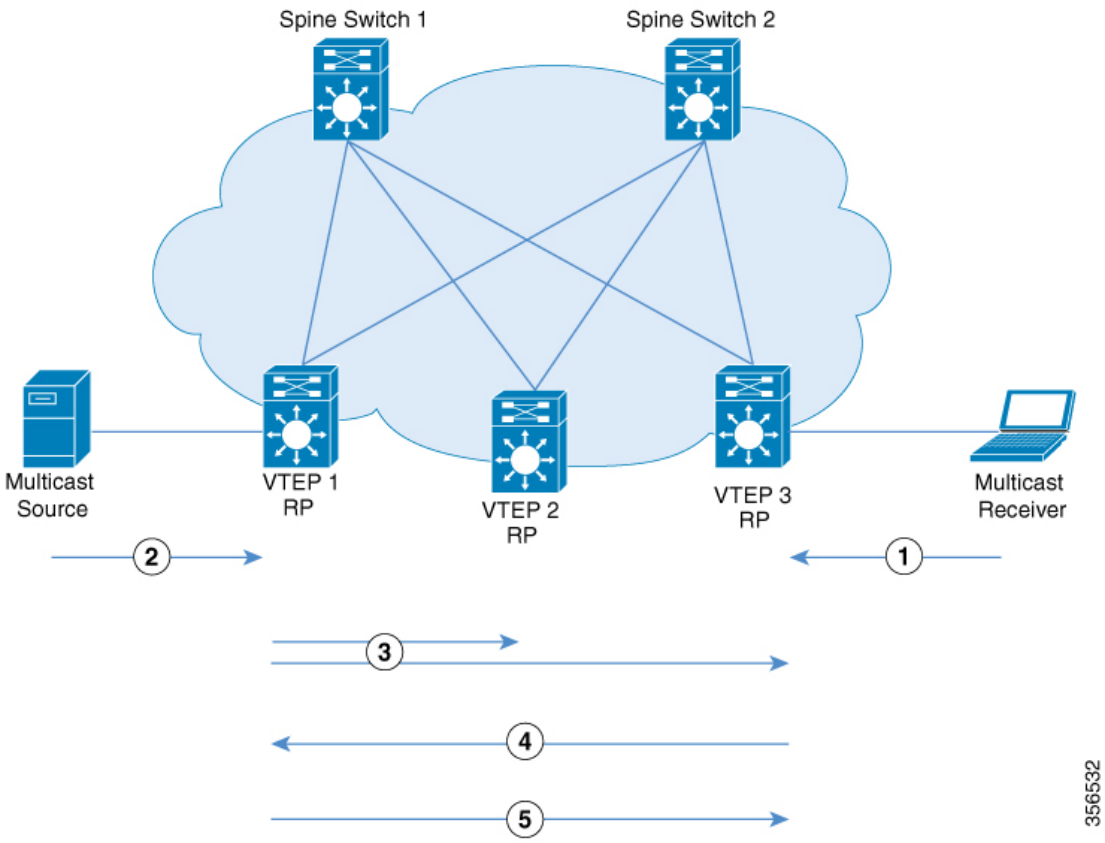
- PIM Sparse Mode with Anycast RP
- PIM Sparse Mode with RP Inside the BGP EVPN VXLAN Fabric
- PIM Sparse Mode with RP Outside the BGP EVPN VXLAN Fabric

PIM Sparse Mode with Anycast RP

In PIM sparse mode with anycast RP, every VTEP in the EVPN VXLAN network acts as an RP in the overlay network for its respective multicast group. The RPs in the underlay network must be configured on the spine switches.

When a VTEP discovers a source device, it sends Source A-D Routes (MVPN route type 5) to all the other VTEPs. Based on these Source A-D routes, the other VTEPs send (S,G) join requests as MVPN route type 7 to the source VTEP.

Figure 38: PIM Sparse Mode with Anycast RP



In PIM sparse mode with anycast RP, the following sequence of events occurs:

1. Receiver sends (*,G) IGMP Join to VTEP 3. Since VTEP 3 is an RP, (*,G) is created at VTEP 3.
2. The source device starts streaming data and (S,G) is created on VTEP 1.
3. VTEP 1 performs self-source-registration since it is also an RP.

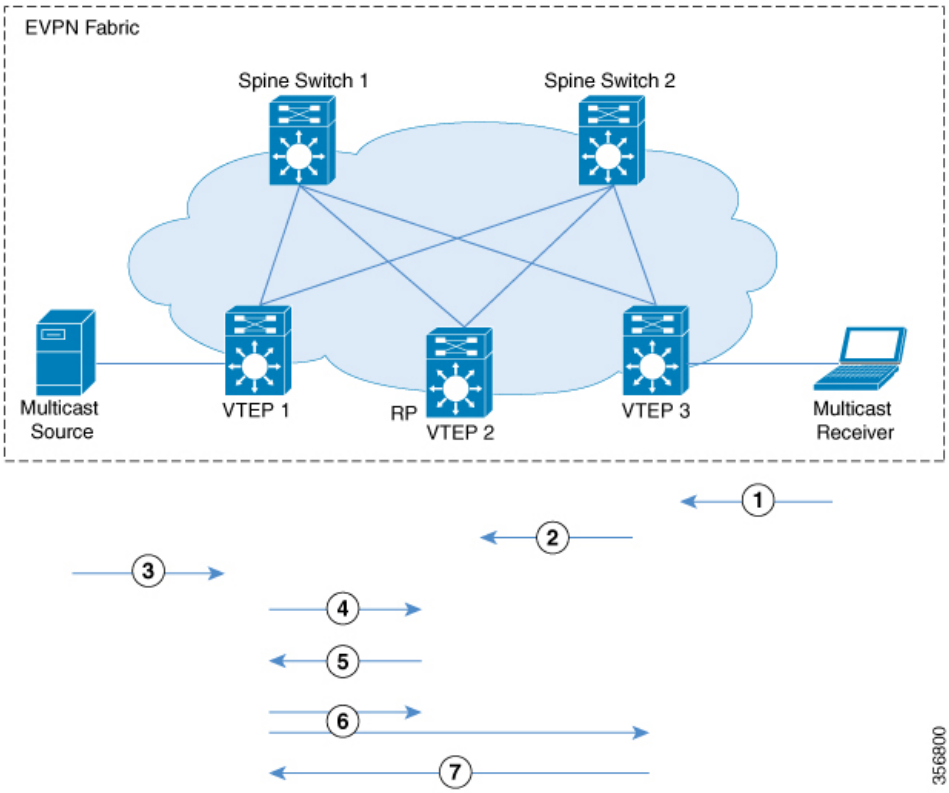
   The source VTEP (VTEP 1) advertises Source A-D Routes (also called MVPN route type 5) for the (S,G) to all the other VTEPs which are BGP peers in the MVPN address family.
4. VTEP 2 and VTEP 3 receive and install the Source A-D Route for the (S,G).

   (S,G) is created at VTEP 3. VTEP 3 now has an overlay route for the (S,G) and also has a unicast route to the source device from the EVPN Control plane. It then sends an MVPN route type 7 (S,G) BGP join to VTEP 1 and starts accepting traffic.
5. VTEP 1 receives and installs MVPN route type 7 from VTEP 3. It uses the Layer 3 VNI's SVI as the forwarding interface for the (S,G) and starts forwarding traffic.

PIM Sparse Mode with RP Inside the BGP EVPN VXLAN Fabric

In PIM sparse mode with RP inside the BGP EVPN VXLAN Fabric, the RP can be any VTEP in the EVPN VXLAN network.

Figure 39: PIM Sparse Mode with RP Inside the BGP EVPN VXLAN Fabric



The following sequence of events occurs when TRM is enabled in PIM sparse mode with the RP inside the fabric:

1. Receiver sends (*,G) IGMP Join to VTEP 3. (*,G) is created at VTEP 3.
2. VTEP 3 sends MVPN route type 6 to VTEP 2 which is the RP. (*,G) is created at VTEP 2.
3. The source device starts streaming data and (S,G) is created on VTEP 1.
4. VTEP 1 performs source-registration at VTEP 2 since it is the RP. (S,G) is created at VTEP 2.
5. Since the RP has a receiver for (S,G), it sends an MVPN route type 7 to VTEP 1 and forwards PIM-register traffic towards receivers in the (*,G) tree.
6. VTEP 1 receives and installs MVPN route type 7 from VTEP 2. It uses the Layer 3 VNI's SVI as the forwarding interface for (S,G).

   The source VTEP (VTEP 1) advertises Source A-D Routes for (S,G) to all the other VTEPs which are BGP peers in the MVPN address family.

   VTEP 2 and VTEP 3 receive and install the Source A-D Routes for (S,G).
7. (S,G) is created at VTEP 3. VTEP 3 now has an overlay route for (S,G) and also has a unicast route to the source device from the EVPN Control plane. It then sends an MVPN route type 7 to VTEP 1 and starts accepting traffic.

   VTEP 1 receives and installs MVPN route type 7 from VTEP 3 and starts forwarding traffic.

Note For the receiver VTEP to be able to send an MVPN route type 7 to the source VTEP, there can be two triggers:

- The (*,G) packets being forwarded to the receiver VTEP from the RP.
- The Source A-D route received from the source VTEP.

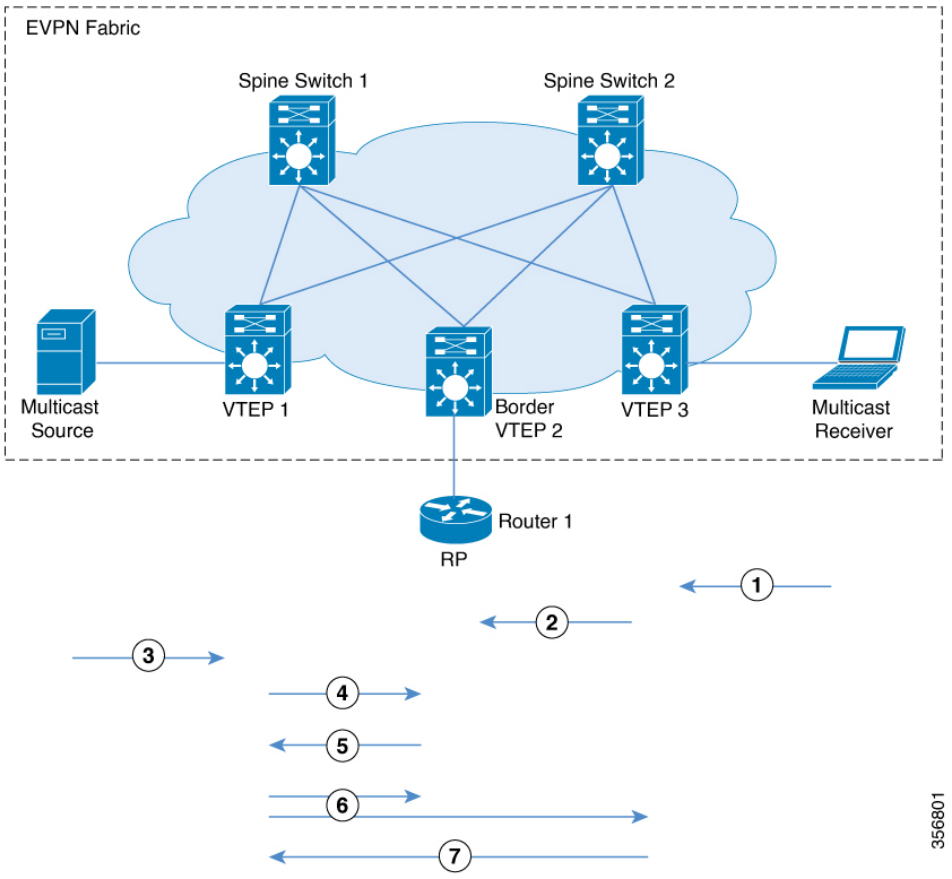
Once either of these are received, the receiver VTEP sends MVPN route type 7 to the source VTEP.

PIM Sparse Mode with RP Outside the BGP EVPN VXLAN Fabric

In PIM sparse mode with RP outside the BGP EVPN VXLAN Fabric, the RP can be a PIM router behind any VTEP in the EVPN VXLAN network.

Note When the RP is configured outside the BGP EVPN VXLAN fabric, TRM in PIM sparse mode functions the same way as it does when the RP is inside the fabric.

Figure 40: PIM Sparse Mode with RP Outside the BGP EVPN VXLAN Fabric



The chronological traffic flow from the image above is as follows:

1. Receiver sends (*,G) IGMP Join to VTEP 3. (*,G) is created at VTEP 3.
2. VTEP3 sends and MVPN route type 6 to VTEP 2 which has the RP in its overlay network. This route is converted to a (*,G) join towards the RP by VTEP2.
3. The source device starts streaming data and (S,G) is created on VTEP 1.
4. VTEP1 performs source registration with RP. (S,G) join from RP creates (S,G) state at VTEP 2.
5. Since the RP has a receiver for (S,G), it sends an MVPN route type 7 to VTEP 1 and forwards PIM-register traffic towards receivers in the (*,G) tree.
6. VTEP 1 receives and installs MVPN route type 7 from VTEP 2. It uses the Layer 3 VNI's SVI as the forwarding interface for (S,G).

   The source VTEP (VTEP 1) advertises Source A-D Routes for (S,G) to all the other VTEPs which are BGP peers in the MVPN address family.

   VTEP 2 and VTEP 3 receive and install the Source A-D Routes for (S,G).
7. (S,G) is created at VTEP 3. VTEP 3 now has an overlay route for (S,G) and also has a unicast route to the source device from the EVPN Control plane. It then sends an MVPN route type 7 to VTEP 1 and starts accepting traffic.

   VTEP 1 receives and installs MVPN route type 7 from VTEP 3 and starts forwarding traffic.

Note For the receiver VTEP to be able to send an MVPN route type 7 to the source VTEP, there can be two triggers:

- The (*,G) packets being forwarded to the receiver VTEP from the RP.
- The Source A-D route received from the source VTEP.

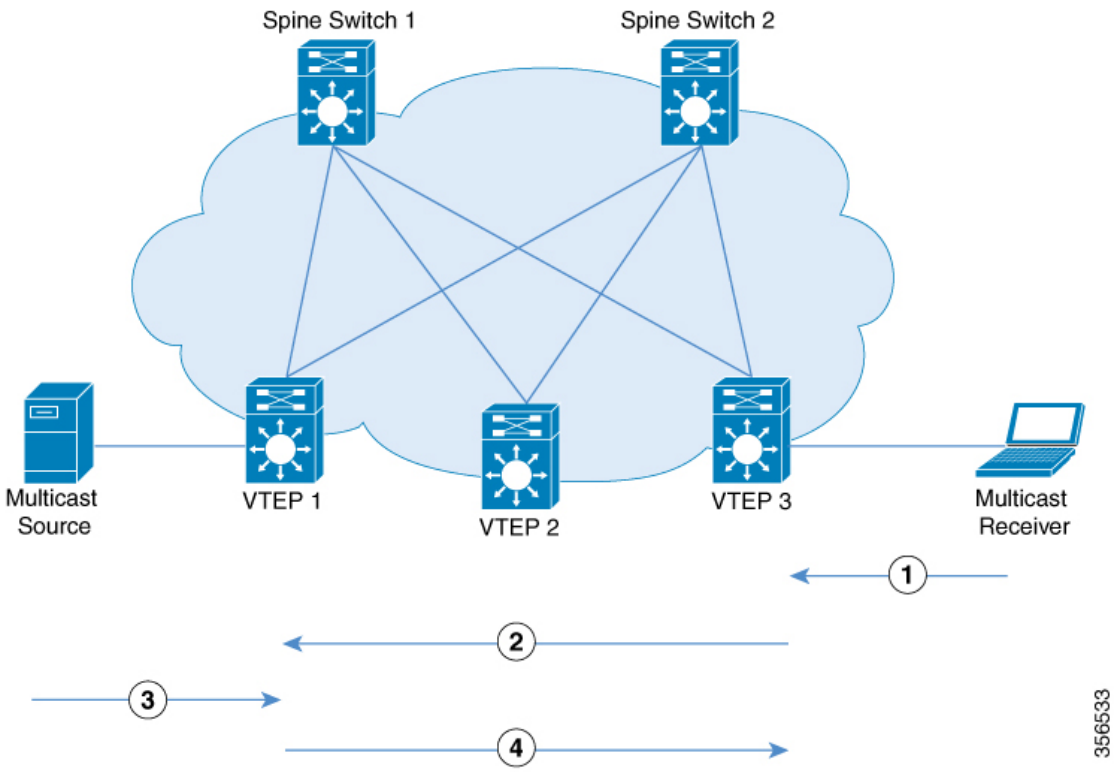
Once either of these are received, the receiver VTEP sends MVPN route type 7 to the source VTEP.

See PIM Sparse mode with RP located inside the BGP EVPN VXLAN fabric for the sequence of events that happen when TRM is enabled in PIM sparse mode with the RP inside the fabric.

TRM in PIM Source Specific Mode

In PIM source specific mode, the Source A-D route (MVPN route type 5) is not needed for the multicast convergence to happen. The receiver VTEP does not wait to receive the Source A-D route to send the MVPN route type 7.

Figure 41: PIM Source Specific Mode



In PIM Source Specific Mode, the following sequence of events occurs:

1. When the source device sends a unicast packet, VTEP 1 sends out EVPN routes to all the other VTEPs, letting them know that the packet is from the source device.

   The receiver sends an (S,G) IGMP join towards VTEP 3 and an (S,G) entry is created.

2. VTEP 3 performs an RPF lookup for the source device. If the SVI of the Layer 3 VNI is found to be the RPF interface, VTEP 3 sends MVPN route type 7 towards VTEP 1.

3. VTEP 1 receives and installs the MVPN route type 7. VTEP 1 creates an (S,G) entry, using the Layer 3 VNI's SVI as the forwarding interface for (S,G).

   The source device sends (S,G) data to VTEP 1.

4. VTEP 1 starts forwarding the traffic to VTEP 3.

Data MDT

Data MDTs are purpose built underlay MDTs to provide optimized forwarding in the MVPN and EVPN core. Threshold and access control configurations can be used to control the characteristics of overlay streams for which data MDTs will be built. The threshold at which the data MDT is created can be configured on a per-VRF basis only. When the multicast transmission exceeds the defined threshold, the sending VTEP device creates the data MDT and sends a MVPN route type 3 message, which contains information about the data MDT, to all devices on the default MDT. The statistics to determine whether a multicast stream has exceeded

the data MDT threshold are examined periodically. After a VTEP device sends the MVPN route type 3 message, it waits 3 more seconds before switching over.

Benefits of Data MDT

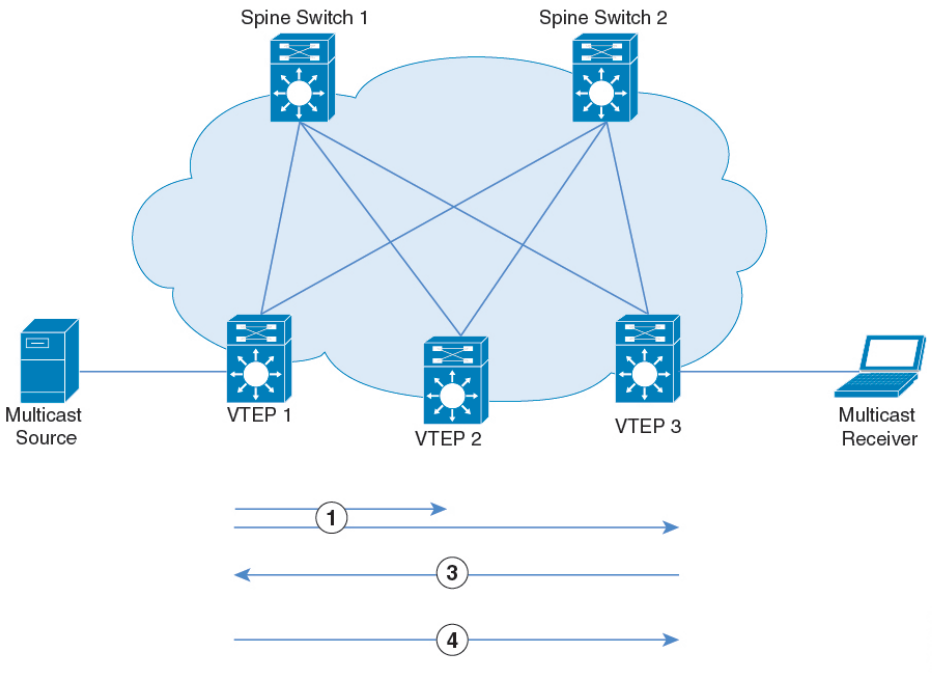
- With Data MDT, replication load on the EVPN spine nodes are greatly reduced and there are less number of copies in the EVPN underlay overall.
- Data MDT lowers the bandwidth usage between the spine nodes and VTEPs. The traffic load at VTEPs that do not require certain overlay streams are completely avoided.
- Data MDT also offers flexibility in deployment through ACLs and threshold-based switchovers, and also enables load-balancing capabilities.

Data MDTs are created only for (S, G) multicast route entries within the VRF multicast routing table. They are not created for (*, G) entries regardless of the value of the individual source data rate. An ACL can be configured to control the overlay streams that will be allowed onto the data MDT irrespective of the threshold.

Data MDT is supported for all TRM modes: PIM sparse mode with anycast RP, PIM sparse mode with RP inside the BGP EVPN VXLAN fabric, PIM sparse mode with RP outside the BGP EVPN VXLAN fabric, and PIM source specific multicast (SSM) mode.

Note For an optimal usage of Data MDT, we recommend that you configure PIM-SSM in the underlay. If you configure PIM sparse mode in the underlay, use distinct Data MDT ranges on each VTEP.

Figure 42: Data MDT Mode



After the sequence of events of PIM sparse mode with RP outside the BGP EVPN VXLAN fabric or PIM source specific mode is completed, continue with the following sequence of events for Data MDT mode:

1. VTEP1 sends MVPN route type 3 to all other nodes. The threshold is exceeded for stream and MVPN route type 3 is sent to advertise data MDT for the stream. After 3 seconds, VTEP1 switches stream from the default MDT to data MDT using MVPN route type 3. Only VTEP3 receives stream data.
2. VTEP2 has no receiver, and so it does not join the advertised data MDT underlay.
3. VTEP3 has a receiver, and joins the advertised data MDT underlay.

How to Configure Tenant Routed Multicast

- Configuring Tenant Routed Multicast (TRM), on page 373

  Configuring TRM using CLI provides concise steps to configure TRM.

- Configuring Tenant Routed Multicast (TRM) Advanced Mode, on page 375

  Configuring TRM using advanced mode provides detail steps with extra customisation.

To configure TRM, you can use any one of the above mentioned methods.

Prerequisites to Configuring TRM

Before configuring TRM, ensure that EVPN VXLAN Layer 2 and Layer 3 Overlay networks have been configured. See How to Configure EVPN VXLAN Integrated Routing and Bridging, on page 108 for detailed steps to configure Layer 2 and Layer 3 overlay networks.

Configuring Tenant Routed Multicast (TRM)

For quick TRM MDT setup, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf rd-auto** **Example:** `Device# vrf rd-auto`	Configures auto route distinguisher (RD) for all the VRFs.
Step 4	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition green`	Names the VRF and enters VRF configuration mode.

	Command or Action	Purpose
Step 5	**vnid** *vnid-number* **evpn-instance vni** *l3vni-number* **core vlan** *core vlan number* **Example:** `Device(config-vrf)# vnid 30500 evpn-instance vni 50000 core-vlan 500`	Creates core VLAN and VNIs for the VRF.
Step 6	**address-family** {**ipv4** | **ipv6** } **Example:** `Device(config-vrf)# address-family ipv4`	Specifies the VRF and enters VRF address family configuration mode. • Use the **ipv4** keyword to configure IPv4 address family. • Use the **ipv6** keyword to configure IPv6 address family.
Step 7	**evpn-mcast** *group-address* **data** *group-address* **anycast** *group-address* **Example:** `Device(config-vrf-af)# evpn-mcast 225.2.2.2 data 232.0.0.0 0.0.0.128 anycast 5.5.5.5`	Creates a default MDT address. Data MDT group and anycast RP are optional parameters. • Use the **data** option to setup Data mdt. • Use the **anycast** option to setup TRM in Distributed Anycast-RP (spt-only) mode. Please refer to the corresponding section in Configuring Tenant Routed Multicast (TRM) Advanced Mode, on page 375
Step 8	**exit-address-family** **Example:** `Device(config-vrf-af)# exit-address-family`	Exits VRF address family configuration mode and returns to configuration mode.
Step 9	**mdt data threshold** *kb/s* **Example:** `Device(config-vrf-af)# mdt data threshold 111`	(Optional) Defines the bandwidth threshold value in kilobits per second (kb/s). The range is from 1 to 4294967, and the default value is 0.
Step 10	**ip pim vrf** *vrf-name* **register-source** *source-address* **Example:** `Device(config)# ip pim vrf BLUE register-source Lo11`	Configures IPv4 PIM register source for the VRF in case of external RP
Step 11	**ipv6 pim vrf** *vrf-name* **register-source** *source-address* **Example:** `Device(config)# ipv6 pim vrf BLUE register-source Lo11`	Configures IPv6 PIM register source for the VRF in case of external RP

	Command or Action	Purpose
Step 12	**ip pim vrf** *vrf-name* **rp-address** *rp-address* **Example:** `Device(config)# ip pim vrf BLUE register-source Lo11`	Configures IPv4 PIM RP address for the VRF in case of external RP
Step 13	**ipv6 pim vrf** *vrf-name* **rp-address** *rp-address* **Example:** `Device(config)# ipv6 pim vrf BLUE register-source Lo11`	Configures IPv6 PIM RP address for the VRF in case of external RP
Step 14	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuring Tenant Routed Multicast (TRM) Advanced Mode

Use this method if you wish to customise TRM configuration, for configuration without customisation refer to Configuring Tenant Routed Multicast (TRM), on page 373.

Perform the following set of procedures to configure TRM in an EVPN VXLAN network with customisation:

Configuring TRM with PIM Sparse Mode

To configure TRM with PIM Sparse Mode, perform the following tasks:

Configuring the TRM Multicast Distribution Tree in the VRF

To configure the TRM MDT, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition green`	Names the VRF and enters VRF configuration mode.

	Command or Action	Purpose
Step 4	**address-family** {**ipv4** | **ipv6** } **Example:** `Device(config-vrf)# address-family ipv4`	Specifies the VRF and enters VRF address family configuration mode. • Use the **ipv4** keyword to configure IPv4 address family. • Use the **ipv6** keyword to configure IPv6 address family.
Step 5	**mdt default vxlan** *group-address* **Example:** `Device(config-vrf-af)# mdt default vxlan 225.2.2.2`	Configures the multicast group address range for default MDT groups for a VRF in a VXLAN.
Step 6	**mdt auto-discovery vxlan** [**inter-as**] **Example:** `Device(config-vrf-af)# mdt auto-discovery vxlan`	Enables VXLAN with BGP auto-discovery. Use the **inter-as** keyword for the MVPN address family routes to cross the BGP autonomous system (AS) boundaries.
Step 7	**mdt overlay use-bgp** [**spt-only**] **Example:** `Device(config-vrf-af)# mdt overlay use-bgp spt-only`	Configures the mechanism that is used by TRM in PIM sparse mode to operate within the BGP EVPN VXLAN fabric. Specifies BGP as the overlay protocol. • Use the **mdt overlay use-bgp spt-only** command to configure PIM sparse mode with anycast RP. • Use the **mdt overlay use-bgp** command to configure PIM sparse mode with a single RP either inside or outside the BGP EVPN VXLAN fabric.
Step 8	**mdt data vxlan** *data-mdt-subnet data-mdt-mask* [**list** *access-list-number*] **Example:** `Device(config-vrf-af)# mdt data vxlan 225.2.2.0 0.0.0.255 list 101`	(Optional) Configures the multicast group address range for data MDT groups for a VRF in a VXLAN.
Step 9	**mdt data threshold** *kb/s* **Example:** `Device(config-vrf-af)# mdt data threshold 111`	(Optional) Defines the bandwidth threshold value in kilobits per second (kb/s). The range is from 1 to 4294967, and the default value is 0.
Step 10	**exit-address-family** **Example:** `Device(config-vrf-af)# exit-address-family`	Exits VRF address family configuration mode and returns to VRF configuration mode.

	Command or Action	Purpose
Step 11	**end** **Example:** `Device(config-vrf)# end`	Returns to privileged EXEC mode.

Configuring Multicast Routing on the Overlay VRF

To enable multicast routing on the overlay VRF, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip multicast-routing vrf** *vrf-name* **Example:** `Device(config)# ip multicast-routing vrf green`	Enables IP multicast forwarding on the overlay VRF.
Step 4	**ipv6 unicast-routing** **Example:** `Device(config)# ipv6 unicast-routing`	Enables IPv6 unicast forwarding.
Step 5	**ipv6 multicast-routing vrf** *vrf-name* **Example:** `Device(config)# ipv6 multicast-routing vrf green`	Enables IPv6 multicast forwarding on the overlay VRF.
Step 6	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuring Multicast on Switch Virtual Interfaces for Core-facing and Access-facing VLANs

To configure multicast on SVIs for the core-facing and access-facing VLANs on the VTEP, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *core-facing-vlan-id* **Example:** `Device(config)# interface vlan 200`	Enters interface configuration mode for the specified VLAN.
Step 4	**ip pim sparse-mode** **Example:** `Device(config-if) # ip pim sparse-mode`	Enables IPv4 multicast on the core-facing SVI.
Step 5	**exit** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.
Step 6	**interface vlan** *access-facing-vlan-id* **Example:** `Device(config)# interface vlan 202`	Enters interface configuration mode for the specified VLAN.
Step 7	**ip pim sparse-mode** **Example:** `Device(config-if) # ip pim sparse-mode`	Enables IPv4 multicast on the access-facing SVI where sources or receivers are connected. Repeat this step for all the access-facing SVIs that are part of the Layer 2 VNI where sources and receivers are connected.
Step 8	**end** **Example:** `Device(config-if)# end`	Returns to privileged EXEC mode.

Configuring BGP with MVPN Address Family on VTEP

To configure BGP on a VTEP with MVPN address family, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable**	Enables privileged EXEC mode.

	Command or Action	Purpose
	Example: `Device> enable`	Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 1`	Enables a BGP routing process, assigns it an autonomous system number, and enters router configuration mode.
Step 4	**address-family** {**ipv4** | **ipv6** } **mvpn** **Example:** `Device(config-router)# address-family ipv4 mvpn`	Specifies the MVPN address family and enters address family configuration mode. • Use the **ipv4** keyword to configure IPv4 MVPN address family. • Use the **ipv6** keyword to configure IPv6 MVPN address family.
Step 5	**neighbor** *ip-address* **activate** **Example:** `Device(config-router-af)# neighbor 10.2.2.20 activate`	Enables the exchange of information with a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 6	**neighbor** *ip-address* **send-community extended** **Example:** `Device(config-router-af)# neighbor 10.2.2.20 send-community both`	Specifies the communities attribute sent to a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 7	**neighbor** *ip-address* **advertisement-interval** *seconds* **Example:** `Device(config-router-af)# neighbor 10.2.2.20 advertisement-interval 10`	(Optional) Sets the minimum route advertisement interval (MRAI) between the sending of BGP routing updates.
Step 8	**exit-address-family** **Example:** `Device(config-router-af)# exit-address-family`	Exits address family configuration mode and returns to router configuration mode.
Step 9	**end** **Example:** `Device(config-router)# end`	Returns to privileged EXEC mode.

Configuring RP for Underlay Network

To configure RP for the underlay network, perform the following steps:



Note We recommend that you configure the Spine Switch as the RP for the underlay network.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip pim rp-address** *ip-address-of-rp* **Example:** `Device(config)# ip pim rp-address <rp-ip-address>`	Configures the RP in the underlay network. For information about RP redundancy, see *IP Multicast Routing Configuration Guide*.
Step 4	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuring RP for Overlay Network

To configure RP for the overlay network, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *loopback-interface* **Example:**	Enters interface configuration mode for the specified Loopback interface.

	Command or Action	Purpose
	`Device(config)# interface Loopback 13`	
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Configures forwarding table for the Loopback interface.
Step 5	**ip-address** *ip-address subnet-mask* **Example:** `Device(config-if)# ip address 10.1.13.13 255.255.255.255`	Configures the IP address for the Loopback interface.
Step 6	**ip pim sparse-mode** **Example:** `Device(config-if)# ip pim sparse-mode`	Enables IPv4 multicast on the Loopback interface. **Note** Enable PIM sparse mode only if EVPN VXLAN Layer 2 overlay network is also configured on the VTEP with underlay multicast as the mechanism for forwarding BUM traffic.
Step 7	**exit** **Example:** `Device(config-if)# exit`	Returns to global configuration mode.
Step 8	{**ip** | **ipv6** } **pim vrf** *vrf-name* **rp-address** *rp-address* **Example:** `Device(config)# ip pim vrf green rp-address 10.1.13.13`	Configures the address of the local VTEP as the PIM RP for the multicast group. • In PIM-SM with Anycast RP mode, use the address of the loopback interface of the local VTEP. • In PIM-SM with RP either inside or outside the BGP EVPN VXLAN fabric, use the IP address of the RP. **Note** The loopback interface specified must be part of the same VRF.
Step 9	{**ip** | **ipv6** } **pim vrf** *vrf-name* **register-source** *loopback-address-of-vtep* **Example:** `Device(config)# ip pim vrf green register-source loopback901`	Configures a unique IP address for the loopback interface of the VTEP that acts as the first hop router to multicast traffic.
Step 10	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Configuring TRM with PIM Source Specific Mode

To configure TRM with PIM Source Specific Mode, perform the following tasks:

- Configure the TRM Multicast Distribution Tree in the VRF
- Configure Multicast Routing on the Overlay VRF
- Configure Multicast on Switch Virtual Interfaces for the Core-facing and Access-facing VLANs
- Configure BGP with MVPN Address Family on VTEP
- Configuring RP for Underlay Network

Configuring SSM for Overlay Network

To configure SSM for the overlay network, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip pim vrf** *vrf-name* **ssm** { **default** | **range** *access-list* } **Example:** `Device(config)# ip pim vrf green ssm default`	Configures an SSM range for TRM. The **default** keyword defines the SSM range access list as 232/8. The **range** keyword specifies the standard IP access list number or name that defines the SSM range.
Step 4	**end** **Example:** `Device(config)# end`	Returns to privileged EXEC mode.

Verifying Tenant Routed Multicast

The following table lists the **show** commands that are used to verify TRM:

Command	Purpose
show nve peers	Displays NVE interface state information for peer leaf switches.

Command	Purpose
show l2vpn evpn peers vxlan	Displays Layer 2 EVPN peer route counts in the VXLAN and up time.
show ip igmp vrf green groups	Displays the multicast groups with receivers that are directly connected to the router pertaining to the specific Multicast Virtual Routing and Forwarding (MVRF) instance and that were learned through IGMP.
show bgp ipv4 mvpn all	Displays the MVPN options for BGP MVPN C-route signaling.
show ip mroute vrf green	Displays the contents of the mroute table that pertain to a specific MVRF instance.
show ip mfib vrf green	Displays forwarding entries and interfaces in the IPv4 Multicast Forwarding Information Base (MFIB) associated with MVRF instances.
show ip mroute	Displays multicast routing table information.
show ip mfib	Displays the forwarding entries and interfaces in the IPv4 MFIB.

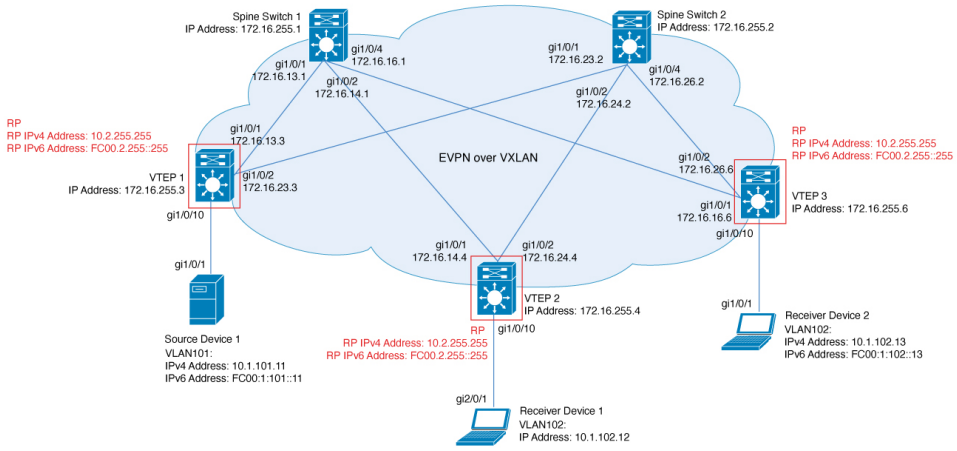
Configuration Examples for Tenant Routed Multicast

The following sections provide configuration examples for TRM in different scenarios.

Note These configuration examples do not have the Data MDT feature enabled.

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

This example shows how to configure and verify Layer 3 TRM with PIM-SM for IPv4 and IPv6 multicast traffic when the every VTEP is an RP is inside the BGP EVPN VXLAN fabric.



The topology shows an EVPN VXLAN network with a receiver device and a source device connected to VTEP 1 and VTEP 2 respectively. The IPv4 multicast group is 226.1.1.1 and the IPv6 multicast group is FF06:1::1 in this topology. The following tables provide sample configurations for the devices in this topology:

Table 39: Configuring VTEP 1, VTEP 2, and VTEP 3 to Configure TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic

VTEP 1	VTEP 2	VTEP 3
Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp spt-only route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp spt-only route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !	Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp spt-only route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp spt-only route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !	Leaf-03# **show running-config** hostname Leaf-03 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp spt-only route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp spt-only route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !

VTEP 1	VTEP 2	VTEP 3
interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.2.255.255 255.255.255.255 ip pim sparse-mode ipv6 address FC00:2:255::255/128 ipv6 enable ! interface Loopback901 vrf forwarding green ip address 10.1.255.1 255.255.255.255 ip pim sparse-mode ipv6 address FC00:1:255::1/128 ipv6 enable ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access !	interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.2.255.255 255.255.255.255 ip pim sparse-mode ipv6 address FC00:2:255::2/128 ipv6 enable ! interface Loopback901 vrf forwarding green ip address 10.1.255.2 255.255.255.255 ip pim sparse-mode ipv6 address FC00:1:255::2/128 ipv6 enable ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 102 switchport mode access !	interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.2.255.255 255.255.255.255 ip pim sparse-mode ipv6 address FC00:2:255::255/128 ipv6 enable ! interface Loopback901 vrf forwarding green ip address 10.1.255.3 255.255.255.255 ip pim sparse-mode ipv6 address FC00:1:255::3/128 ipv6 enable ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 102 switchport mode access !

VTEP 1	VTEP 2	VTEP 3
interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:102::1/64 ipv6 enable ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 !	interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:102::1/64 ipv6 enable ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 !	interface Vlan101 vrf forwarding green ip pim sparse-mode ipv6 address FC00:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:102::1/64 ipv6 enable ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 !

VTEP 1	VTEP 2	VTEP 3
address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end Leaf-01#	address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end Leaf-02#	address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end Leaf-03#

Table 40: Configuring Spine Switch 1 and Spine Switch 2 to Configure TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic

Spine Switch 1	Spine Switch 2
Spine-01# **show running-config** hostname Spine-01 ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.1 255.255.255.255 ip ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/4 no switchport ip address 172.16.16.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 neighbor 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 !	Spine-02# **show running-config** hostname Spine-02 ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.2 255.255.255.255 ip ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/4 no switchport ip address 172.16.26.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp router-id 172.16.255.2 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 neighbor 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 !

Spine Switch 1	Spine Switch 2
address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end Spine-01#	address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end Spine-02#

Verifying TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic

The following sections provide sample outputs for **show** commands to verify TRM with PIM-SM on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 391
- Outputs to Verify the Configuration on VTEP 2, on page 398
- Outputs to Verify the Configuration on VTEP 3, on page 405
- Outputs to Verify the Configuration on Spine Switch 1, on page 413

• Outputs to Verify the Configuration on Spine Switch 2, on page 417

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Leaf-01# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 01:47:43
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 01:47:43
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 01:47:43
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 01:47:43
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  01:47:43
nve1       10102    L2CP 172.16.254.6     7              10102      UP   N/A  01:47:43
Leaf-01#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
Leaf-01# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                  Num routes eVNI     UP time
--------- -------- -----------------------  ---------- -------- --------
nve1      10102    172.16.254.4             7          10102    01:47:43
nve1      10102    172.16.254.6             7          10102    01:47:43


Leaf-01#show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 8, main routing table version 8
3 network entries using 1176 bytes of memory
4 path entries using 640 bytes of memory
2/2 BGP path/bestpath attribute entries using 608 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
14 BGP extended community entries using 1848 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4456 total bytes of memory
BGP activity 69/0 prefixes, 92/2 paths, scan interval 60 secs
3 networks peaked at 11:32:31 Sep 16 2020 UTC (01:42:43.716 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     140     127        8    0    0 01:48:48        1
172.16.255.2    4        65001     146     128        8    0    0 01:48:40        1
Leaf-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 1:

```
Leaf-01# show bgp ipv6 mvpn all
BGP table version is 8, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>   [5][1:1][FC00:1:101::11][FF06:1::1]/42
```

```
                      ::                                    32768 ?
 *>i  [7][1:1][65001][FC00:1:101::11][FF06:1::1]/46
                      172.16.255.4            0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46
                      172.16.255.4            0    100      0 ?
 *>i                   172.16.255.4            0    100      0 ?
Leaf-01#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 1:

```
Leaf-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 65, main routing table version 65
42 network entries using 16128 bytes of memory
60 path entries using 12720 bytes of memory
11/11 BGP path/bestpath attribute entries using 3168 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
14 BGP extended community entries using 1848 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 34048 total bytes of memory
BGP activity 69/0 prefixes, 92/2 paths, scan interval 60 secs
42 networks peaked at 11:27:30 Sep 16 2020 UTC (01:47:45.010 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     140     127       65    0    0 01:48:48        18
172.16.255.2    4        65001     146     128       65    0    0 01:48:40        18
Leaf-01#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 1:

```
Leaf-01# show bgp l2vpn evpn summary
BGP table version is 65, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>   [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                      ::                                    32768 ?
Route Distinguisher: 172.16.254.3:102
 *>i  [2][172.16.254.3:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                      172.16.254.6            0    100      0 ?
```

```
 *>i  [2][172.16.254.3:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][ECE1A93792C5][0][*]/20
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.3:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                     172.16.254.6             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.6:102
 *>i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][0][*]/20
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
Route Distinguisher: 1:1 (default for vrf green)
 *>   [5][1:1][0][24][10.1.101.0]/17
                     0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][24][10.1.102.0]/17
```

```
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>   [5][1:1][0][32][10.1.255.1]/17
                     0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][32][10.1.255.2]/17
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [5][1:1][0][32][10.1.255.3]/17
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
 *>   [5][1:1][0][32][10.2.255.255]/17
                     0.0.0.0                  0         32768 ?
 *>   [5][1:1][0][64][FC00:1:101::]/29
                     ::                       0         32768 ?
 *>i  [5][1:1][0][64][FC00:1:102::]/29
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>   [5][1:1][0][128][FC00:1:255::1]/29
                     ::                       0         32768 ?
 *>i  [5][1:1][0][128][FC00:1:255::2]/29
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [5][1:1][0][128][FC00:1:255::3]/29
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
 *>   [5][1:1][0][128][FC00:2:255::255]/29
                     ::                       0         32768 ?
Leaf-01#
```

The following example shows the output for the **show ipv6 pim vrf** *vrf-name* **group-map** command on VTEP 1:

```
Leaf-01# show ipv6 pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Tu2,FC00:2:255::255 (us)
    Info source: Static
    Uptime: 01:49:05, Groups: 1
Leaf-01#
```

The following example shows the output for the **show ipv6 route vrf** command on VTEP 1:

```
Leaf-01# show ipv6 route vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "connected", distance 0, metric 0, type receive, connected
  Redistributing via bgp 65001
  Route count is 1/1, share count 0
  Routing paths:
    receive via Loopback255
      Last updated 01:49:06 ago
Leaf-01#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 1:

```
Leaf-01# show ipv6 mld vrf green groups
No groups found.
Leaf-01#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(FC00:1:101::11, FF06:1::1), 01:42:44/00:03:19, flags: SFTGq
  Incoming interface: Vlan101
  RPF nbr: FE80::F6CF:E2FF:FE43:34C1
  Immediate Outgoing interface list:
    Vlan901, Forward, 01:42:44/never
Leaf-01#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (FC00:1:101::11,FF06:1::1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 1/0/1
   HW Forwarding:   3161/0/118/0, Other: 0/0/0
   Vlan101 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
Leaf-01#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Leaf-01# show ip mroute
IP Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 01:48:56/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:48:56/00:02:56

(172.16.254.3, 239.1.1.1), 01:42:42/00:03:20, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 01:42:42/00:03:03

(*, 224.0.1.40), 01:49:06/00:02:55, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 01:49:05/00:02:55

(*, 225.0.0.102), 01:48:56/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:48:56/00:02:56

(172.16.254.4, 225.0.0.102), 01:48:17/00:01:19, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:48:17/00:02:56

(172.16.254.6, 225.0.0.102), 01:48:23/00:01:18, flags: Tx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:48:23/00:02:56

(*, 225.0.0.101), 01:49:01/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:48:56/00:02:56

(172.16.254.3, 225.0.0.101), 01:49:01/00:02:45, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 01:48:25/00:03:12
Leaf-01#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Leaf-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
```

```
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 2/0/125/0, Other: 1/0/1
   HW Forwarding:   554/0/163/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   561/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 1/1/0
   HW Forwarding:   504/0/205/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
```

```
 (172.16.254.3,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 1/1/0
   HW Forwarding:   3071/0/156/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
Leaf-01#
```

Return to Verifying TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic, on page 390

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Leaf-02# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 01:52:57
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 01:52:57
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 01:52:57
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 01:52:57
nve1       10101    L2CP 172.16.254.3     7              10101      UP   N/A  01:52:57
nve1       10102    L2CP 172.16.254.6     7              10102      UP   N/A  01:52:57
Leaf-02#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
Leaf-02# show l2vpn evpn peers vxlan

Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            7          10101    01:52:57
nve1      10102    172.16.254.6            7          10102    01:52:57
Leaf-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on VTEP 2:

```
Leaf-02# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 5, main routing table version 5
2 network entries using 784 bytes of memory
3 path entries using 480 bytes of memory
2/2 BGP path/bestpath attribute entries using 608 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
14 BGP extended community entries using 1848 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3904 total bytes of memory
BGP activity 70/0 prefixes, 101/6 paths, scan interval 60 secs
2 networks peaked at 11:37:07 Sep 16 2020 UTC (01:47:58.150 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     150     133        5    0    0 01:53:34        1
172.16.255.2    4        65001     151     134        5    0    0 01:53:30        1
Leaf-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 2:

```
Leaf-02# show bgp ipv6 mvpn all
BGP table version is 5, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][FC00:1:101::11][FF06:1::1]/42
                     172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 *>   [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46
                     ::                                    32768 ?
Leaf-02#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 2:

```
Leaf-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 43, main routing table version 43
42 network entries using 16128 bytes of memory
64 path entries using 13568 bytes of memory
12/12 BGP path/bestpath attribute entries using 3456 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
14 BGP extended community entries using 1848 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 35184 total bytes of memory
BGP activity 70/0 prefixes, 101/6 paths, scan interval 60 secs
42 networks peaked at 11:32:07 Sep 16 2020 UTC (01:52:58.436 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     150     133       43    0    0 01:53:35       20
172.16.255.2    4        65001     151     134       43    0    0 01:53:31       20
Leaf-02#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 2:

```
Leaf-02# show bgp l2vpn evpn
BGP table version is 43, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                     172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
```

```
 *>i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                     172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                     172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                     172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                     172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                     172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [2][172.16.254.4:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][0][*]/20
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                     172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [2][172.16.254.4:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     172.16.254.6             0    100      0 ?
 *>   [2][172.16.254.4:102][0][48][44D3CA286CC5][0][*]/20
                     ::                                    32768 ?
 *>   [2][172.16.254.4:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                     ::                                    32768 ?
 *>   [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                     ::                                    32768 ?
 *>   [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                     ::                                    32768 ?
 *>   [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                     ::                                    32768 ?
 *>   [2][172.16.254.4:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
                     ::                                    32768 ?
 *>i  [2][172.16.254.4:102][0][48][ECE1A93792C5][0][*]/20
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                     172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                     172.16.254.6             0    100      0 ?
Route Distinguisher: 172.16.254.6:102
 *>i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][0][*]/20
                     172.16.254.6             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
```

```
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][0][24][10.1.101.0]/17
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>   [5][1:1][0][24][10.1.102.0]/17
                     0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][32][10.1.255.1]/17
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>   [5][1:1][0][32][10.1.255.2]/17
                     0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][32][10.1.255.3]/17
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][32][10.2.255.255]/17
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>                    0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][64][FC00:1:101::]/29
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>   [5][1:1][0][64][FC00:1:102::]/29
                     ::                       0         32768 ?
 *>i  [5][1:1][0][128][FC00:1:255::1]/29
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>   [5][1:1][0][128][FC00:1:255::2]/29
                     ::                       0         32768 ?
 *>i  [5][1:1][0][128][FC00:1:255::3]/29
                     172.16.254.6             0    100      0 ?
 * i                   172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][128][FC00:2:255::255]/29
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>                    ::                       0         32768 ?
Leaf-02#
```

The following example shows the output for the **show ipv6 pim vrf** *vrf-name* **group-map** command on VTEP 2:

```
Leaf-02# show ipv6 pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Tu2,FC00:2:255::255 (us)
    Info source: Static
    Uptime: 01:54:21, Groups: 1
Leaf-02#
```

The following example shows the output for the **show ipv6 route vrf** command on VTEP 2:

```
Leaf-02# show ipv6 route vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "connected", distance 0, metric 0, type receive, connected
  Redistributing via bgp 65001
  Route count is 1/1, share count 0
  Routing paths:
    receive via Loopback255
      Last updated 01:54:21 ago
Leaf-02#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 2:

```
Leaf-02# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan102
   01:53:45  00:03:52
Leaf-02#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 01:53:45/never, RP FC00:2:255::255, flags: SCJ
  Incoming interface: Tunnel2
  RPF nbr: FC00:2:255::255
  Immediate Outgoing interface list:
    Vlan102, Forward, 01:53:45/never

(FC00:1:101::11, FF06:1::1), 01:47:58/never, flags: STgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.3
  Inherited Outgoing interface list:
    Vlan102, Forward, 01:53:45/never
Leaf-02#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
```

```
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel2 Flags: A NS
   Vlan102 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:101::11,FF06:1::1) Flags: HW
   SW Forwarding: 1/0/100/0, Other: 0/0/0
   HW Forwarding:   3225/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan102 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
Leaf-02#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Leaf-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 01:54:12/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:54:12/00:00:41

(172.16.254.3, 239.1.1.1), 01:47:56/00:02:39, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:47:56/00:00:03

(*, 224.0.1.40), 01:54:21/00:02:39, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 01:54:20/00:02:39
```

```
(*, 225.0.0.102), 01:54:12/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:54:12/00:00:41

(172.16.254.6, 225.0.0.102), 01:53:36/00:00:58, flags: Tx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:53:36/00:00:41

(172.16.254.4, 225.0.0.102), 01:53:47/00:02:53, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 01:53:36/00:02:40, A

(*, 225.0.0.101), 01:54:12/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:54:12/00:00:41

(172.16.254.3, 225.0.0.101), 01:53:11/00:01:39, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:53:11/00:00:48
Leaf-02#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Leaf-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/190/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
```

```
    SW Forwarding: 1/0/172/0, Other: 0/0/0
    HW Forwarding:   529/0/177/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (*,225.0.0.102) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,225.0.0.102) Flags: HW
    SW Forwarding: 2/0/163/0, Other: 3/1/2
    HW Forwarding:   631/0/163/0, Other: 0/0/0
    Null0 Flags: A
    GigabitEthernet1/0/2 Flags: F
      Pkts: 0/0/2    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   530/0/205/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   1/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.3,239.1.1.1) Flags: HW
    SW Forwarding: 1/0/150/0, Other: 0/0/0
    HW Forwarding:   3224/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
Leaf-02#
```

Return to Verifying TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic, on page 390

Outputs to Verify the Configuration on VTEP 3

The following example shows the output for the **show nve peers** command on VTEP 3:

```
Leaf-03# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 02:01:22
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 02:01:22
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 02:01:22
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 02:01:22
nve1       10101    L2CP 172.16.254.3     7              10101      UP   N/A  02:01:22
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  02:01:22
Leaf-03#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 3:

```
Leaf-03# show l2vpn evpn peers vxlan
```

```
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            7          10101    02:01:23
nve1      10102    172.16.254.4            7          10102    02:01:23
Leaf-03#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on VTEP 3:

```
Leaf-03# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.6, local AS number 65001
BGP table version is 5, main routing table version 5
2 network entries using 784 bytes of memory
3 path entries using 480 bytes of memory
2/2 BGP path/bestpath attribute entries using 608 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
14 BGP extended community entries using 1848 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 3904 total bytes of memory
BGP activity 66/0 prefixes, 97/0 paths, scan interval 60 secs
2 networks peaked at 11:29:08 Sep 16 2020 UTC (01:56:22.908 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     160     143        5    0    0 02:01:59        1
172.16.255.2    4        65001     159     142        5    0    0 02:01:59        1
Leaf-03#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 3:

```
Leaf-03# show bgp ipv6 mvpn all
BGP table version is 5, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][FC00:1:101::11][FF06:1::1]/42
                     172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 *>   [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46
                     ::                                   32768 ?
Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on VTEP 3:

```
Leaf-03# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.6, local AS number 65001
BGP table version is 51, main routing table version 51
42 network entries using 16128 bytes of memory
68 path entries using 14416 bytes of memory
12/12 BGP path/bestpath attribute entries using 3456 bytes of memory
```

```
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
14 BGP extended community entries using 1848 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 36032 total bytes of memory
BGP activity 66/0 prefixes, 97/0 paths, scan interval 60 secs
42 networks peaked at 11:24:07 Sep 16 2020 UTC (02:01:24.200 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     160     143       51    0    0 02:02:00        22
172.16.255.2    4        65001     159     142       51    0    0 02:01:59        22
Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 3:

```
Leaf-03# show bgp l2vpn evpn
BGP table version is 51, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][0][*]/20
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.4:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
```

```
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.6:101
 *>i  [2][172.16.254.6:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.6:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.6:101][0][48][F4CFE24334C1][0][*]/20
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.6:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.6:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                     172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.6:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                     172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.6:102
 *>   [2][172.16.254.6:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                     ::                                    32768 ?
 *>   [2][172.16.254.6:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     ::                                    32768 ?
 *>i  [2][172.16.254.6:102][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                     172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.6:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
                     172.16.254.4             0    100      0 ?
 *>   [2][172.16.254.6:102][0][48][ECE1A93792C5][0][*]/20
                     ::                                    32768 ?
 *>   [2][172.16.254.6:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
                     ::                                    32768 ?
 *>   [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                     ::                                    32768 ?
 *>   [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                     ::                                    32768 ?
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][0][24][10.1.101.0]/17
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][24][10.1.102.0]/17
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>                    0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][32][10.1.255.1]/17
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [5][1:1][0][32][10.1.255.2]/17
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>   [5][1:1][0][32][10.1.255.3]/17
                     0.0.0.0                  0         32768 ?
 * i  [5][1:1][0][32][10.2.255.255]/17
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>                    0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][64][FC00:1:101::]/29
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][64][FC00:1:102::]/29
```

```
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>                    ::                       0         32768 ?
 *>i  [5][1:1][0][128][FC00:1:255::1]/29
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i  [5][1:1][0][128][FC00:1:255::2]/29
                     172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>   [5][1:1][0][128][FC00:1:255::3]/29
                     ::                       0         32768 ?
 * i  [5][1:1][0][128][FC00:2:255::255]/29
                     172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>                    ::                       0         32768 ?
Leaf-03#
```

The following example shows the output for the **show ipv6 pim vrf** *vrf-name* **group-map** command on VTEP 3:

```
Leaf-03# show ipv6 pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Tu2,FC00:2:255::255 (us)
    Info source: Static
    Uptime: 02:02:54, Groups: 1
Leaf-03#
```

The following example shows the output for the **show ipv6 route vrf** command on VTEP 3:

```
Leaf-03# show ipv6 route vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "connected", distance 0, metric 0, type receive, connected
  Redistributing via bgp 65001
  Route count is 1/1, share count 0
  Routing paths:
    receive via Loopback255
      Last updated 02:02:55 ago
Leaf-03#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 3:

```
Leaf-03# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan102
   02:02:06  00:03:50
Leaf-03#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 3:

```
Leaf-03# show ipv6 mroute vrf green
Multicast Routing Table
```

```
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 02:02:06/never, RP FC00:2:255::255, flags: SCJ
  Incoming interface: Tunnel2
  RPF nbr: FC00:2:255::255
  Immediate Outgoing interface list:
    Vlan102, Forward, 02:02:06/never

(FC00:1:101::11, FF06:1::1), 01:56:23/never, flags: STgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.3
  Inherited Outgoing interface list:
    Vlan102, Forward, 02:02:06/never
Leaf-03#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 3:

```
Leaf-03# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel2 Flags: A NS
   Vlan102 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:101::11,FF06:1::1) Flags: HW
   SW Forwarding: 1/0/100/0, Other: 0/0/0
   HW Forwarding:   3475/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan102 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
Leaf-03#
```

The following example shows the output for the **show ip mroute** command on VTEP 3:

```
Leaf-03# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 02:02:45/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 02:02:45/00:01:07

(172.16.254.3, 239.1.1.1), 01:56:21/00:02:07, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 01:56:21/00:00:37

(*, 224.0.1.40), 02:02:55/00:02:10, RP 0.0.0.0, flags: DCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback0, Forward/Sparse, 02:02:54/00:02:10

(*, 225.0.0.102), 02:02:45/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 02:02:45/00:01:07

(172.16.254.4, 225.0.0.102), 02:01:56/00:02:46, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 02:01:56/00:01:07

(172.16.254.6, 225.0.0.102), 02:02:08/00:02:37, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 02:01:58/00:02:59

(*, 225.0.0.101), 02:02:45/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 02:02:45/00:01:07

(172.16.254.3, 225.0.0.101), 02:01:36/00:01:06, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 02:01:36/00:01:23
Leaf-03#
```

The following example shows the output for the **show ip mfib** command on VTEP 3:

```
Leaf-03# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel0, VXLAN Decap Flags: NS
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.0/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/190/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/172/0, Other: 0/0/0
   HW Forwarding:   568/0/177/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   632/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 2/0/180/0, Other: 3/0/3
   HW Forwarding:   610/0/189/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
```

```
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   3474/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
Leaf-03#
```

Return to Verifying TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic, on page 390

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Spine Switch 1:

```
Spine-01# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 20, main routing table version 20
2 network entries using 784 bytes of memory
5 path entries using 800 bytes of memory
2/2 BGP path/bestpath attribute entries using 608 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
13 BGP extended community entries using 1808 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4144 total bytes of memory
BGP activity 1001/969 prefixes, 7359/7288 paths, scan interval 60 secs
2 networks peaked at 11:16:15 Sep 16 2020 UTC (02:20:36.059 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001     148     161       20    0    0 02:08:00        1
172.16.255.4    4        65001     148     165       20    0    0 02:07:32        1
172.16.255.6    4        65001     149     166       20    0    0 02:07:32        1
Spine-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Spine Switch 1:

```
Spine-01# show bgp ipv6 mvpn all
BGP table version is 20, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][FC00:1:101::11][FF06:1::1]/42
                     172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46
```

```
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i                  172.16.255.6             0    100      0 ?
Spine-01#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 785, main routing table version 785
30 network entries using 10320 bytes of memory
66 path entries using 13728 bytes of memory
11/11 BGP path/bestpath attribute entries using 3168 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
13 BGP extended community entries using 1808 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 29168 total bytes of memory
BGP activity 1001/969 prefixes, 7359/7288 paths, scan interval 60 secs
44 networks peaked at 10:13:07 Aug 6 2020 UTC (5w6d ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001     148     161      785    0    0 02:08:00        12
172.16.255.4    4        65001     148     165      785    0    0 02:07:33        12
172.16.255.6    4        65001     149     166      785    0    0 02:07:33        12
Spine-01#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn
BGP table version is 785, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
```

```
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.6:102
 * i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][0][*]/20
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 *>i  [5][1:1][0][24][10.1.102.0]/17
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][32][10.1.255.1]/17
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][32][10.1.255.2]/17
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [5][1:1][0][32][10.1.255.3]/17
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][32][10.2.255.255]/17
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.6             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][64][FC00:1:101::]/29
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 *>i  [5][1:1][0][64][FC00:1:102::]/29
                     172.16.254.4             0    100      0 ?
```

```
 * i                    172.16.254.4             0    100      0 ?
 * i                    172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][128][FC00:1:255::1]/29
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][128][FC00:1:255::2]/29
                       172.16.254.4             0    100      0 ?
 *>i                    172.16.254.4             0    100      0 ?
 * i  [5][1:1][0][128][FC00:1:255::3]/29
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][128][FC00:2:255::255]/29
                       172.16.254.4             0    100      0 ?
 * i                    172.16.254.6             0    100      0 ?
 * i                    172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
Spine-01#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 1:

```
Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-01#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 1:

```
Spine-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 224.0.1.40), 1w2d/00:02:38, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1w2d/00:02:38

(*, 225.0.0.102), 1w2d/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.4, 225.0.0.102), 02:04:11/00:02:37, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.14.4
  Outgoing interface list: Null
```

```
(172.16.254.6, 225.0.0.102), 02:07:34/00:00:34, flags: PTA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null
Spine-01#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 1:

```
Spine-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 84/84/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 4/0/193/0, Other: 41/0/41
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: NS
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/206/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/4 Flags: A NS
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
Spine-01#
```

Return to Verifying TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic, on page 390

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Spine Switch 2:

```
Spine-02# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
```

```
BGP table version is 20, main routing table version 20
2 network entries using 784 bytes of memory
5 path entries using 800 bytes of memory
2/2 BGP path/bestpath attribute entries using 608 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
13 BGP extended community entries using 1808 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 4144 total bytes of memory
BGP activity 1031/999 prefixes, 7443/7372 paths, scan interval 60 secs
2 networks peaked at 11:17:12 Sep 16 2020 UTC (02:22:21.833 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001     150     169       20    0    0 02:09:38        1
172.16.255.4    4        65001     151     168       20    0    0 02:09:14        1
172.16.255.6    4        65001     150     167       20    0    0 02:09:18        1
Spine-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Spine Switch 2:

```
Spine-02# show bgp ipv6 mvpn all
BGP table version is 20, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][FC00:1:101::11][FF06:1::1]/42
                      172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 *>i  [7][172.16.254.3:101][65001][FC00:1:101::11][FF06:1::1]/46
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
Spine-02#
```

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 712, main routing table version 712
30 network entries using 10320 bytes of memory
66 path entries using 13728 bytes of memory
11/11 BGP path/bestpath attribute entries using 3168 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
13 BGP extended community entries using 1808 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 29168 total bytes of memory
BGP activity 1031/999 prefixes, 7443/7372 paths, scan interval 60 secs
44 networks peaked at 10:13:54 Aug 6 2020 UTC (5w6d ago)
```

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001     150     169      712    0    0 02:09:38        12
172.16.255.4    4        65001     151     168      712    0    0 02:09:15        12
172.16.255.6    4        65001     150     167      712    0    0 02:09:18        12
Spine-02#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn
BGP table version is 712, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FC00:1:101::1]/36
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FC00:1:101::11]/36
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][32][10.1.102.12]/24
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FC00:1:102::12]/36
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][44D3CA286CC5][128][FE80::46D3:CAFF:FE28:6CC5]/36
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
 * i  [2][172.16.254.4:102][0][48][7C210DBD954D][128][FC00:1:102::1]/36
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.6:102
 * i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][32][10.1.102.1]/24
                     172.16.254.6             0    100      0 ?
 *>i                  172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][0C75BD67EF4D][128][FC00:1:102::1]/36
                     172.16.254.6             0    100      0 ?
```

```
 *>i                    172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][0][*]/20
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][32][10.1.102.13]/24
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FC00:1:102::13]/36
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
 * i  [2][172.16.254.6:102][0][48][ECE1A93792C5][128][FE80::EEE1:A9FF:FE37:92C5]/36
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 *>i  [5][1:1][0][24][10.1.102.0]/17
                       172.16.254.4             0    100      0 ?
 * i                    172.16.254.4             0    100      0 ?
 * i                    172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][32][10.1.255.1]/17
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][32][10.1.255.2]/17
                       172.16.254.4             0    100      0 ?
 *>i                    172.16.254.4             0    100      0 ?
 * i  [5][1:1][0][32][10.1.255.3]/17
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][32][10.2.255.255]/17
                       172.16.254.4             0    100      0 ?
 * i                    172.16.254.6             0    100      0 ?
 * i                    172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][64][FC00:1:101::]/29
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 *>i  [5][1:1][0][64][FC00:1:102::]/29
                       172.16.254.4             0    100      0 ?
 * i                    172.16.254.4             0    100      0 ?
 * i                    172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][128][FC00:1:255::1]/29
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][128][FC00:1:255::2]/29
                       172.16.254.4             0    100      0 ?
 *>i                    172.16.254.4             0    100      0 ?
 * i  [5][1:1][0][128][FC00:1:255::3]/29
                       172.16.254.6             0    100      0 ?
 *>i                    172.16.254.6             0    100      0 ?
 * i  [5][1:1][0][128][FC00:2:255::255]/29
                       172.16.254.4             0    100      0 ?
 * i                    172.16.254.6             0    100      0 ?
 * i                    172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
Spine-02#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 2:

```
Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings
```

```
Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-02#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 2:

```
Spine-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 22:51:54/00:03:26, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 02:09:47/00:03:26
    GigabitEthernet1/0/2, Forward/Sparse, 02:09:20/00:02:34
    GigabitEthernet1/0/4, Forward/Sparse, 02:09:16/00:03:12

(172.16.254.3, 239.1.1.1), 02:03:40/00:02:43, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 02:03:40/00:03:12
    GigabitEthernet1/0/2, Forward/Sparse, 02:03:40/00:02:46

(*, 224.0.1.40), 1w2d/00:03:18, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:02:09/00:03:18
    GigabitEthernet1/0/1, Forward/Sparse, 00:02:10/00:03:17
    Loopback2, Forward/Sparse, 1w2d/00:02:45

(*, 225.0.0.102), 1w2d/00:03:22, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 02:09:47/00:02:35
    GigabitEthernet1/0/2, Forward/Sparse, 02:09:20/00:03:16
    GigabitEthernet1/0/4, Forward/Sparse, 02:09:16/00:03:22

(172.16.254.6, 225.0.0.102), 02:09:47/00:01:33, flags: MT
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 02:09:16/00:03:16
    GigabitEthernet1/0/1, Forward/Sparse, 02:09:16/00:03:05

(172.16.254.4, 225.0.0.102), 02:09:47/00:02:06, flags: MT
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 02:09:19/00:03:10
```

```
    GigabitEthernet1/0/4, Forward/Sparse, 02:09:16/00:03:22

(*, 225.0.0.101), 1w2d/00:03:29, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 02:09:47/00:03:29
    GigabitEthernet1/0/2, Forward/Sparse, 02:09:20/00:02:31
    GigabitEthernet1/0/4, Forward/Sparse, 02:09:16/00:03:29

(172.16.254.3, 225.0.0.101), 02:09:22/00:03:25, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 02:09:20/00:02:39
    GigabitEthernet1/0/4, Forward/Sparse, 02:09:16/00:03:29
Spine-02#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 2:

```
Spine-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 1/1/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 2/0/140/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   604/0/178/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A NS
```

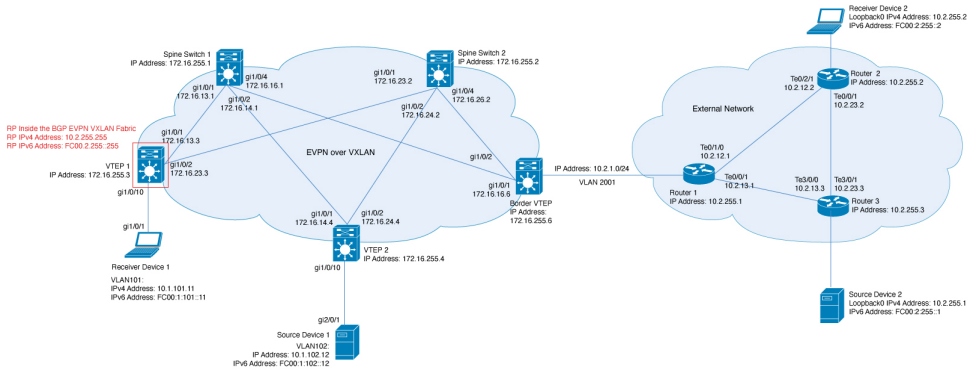
```
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (*,225.0.0.102) Flags: C HW
    SW Forwarding: 4/0/124/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   668/0/176/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   600/0/205/0, Other: 0/0/0
    GigabitEthernet1/0/4 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 17/0/397/0, Other: 4/4/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.3,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   3693/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/1 Flags: A
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
Spine-02#
```

Return to Verifying TRM in PIM-SM with Anycast RP for IPv4 and IPv6 Multicast Traffic, on page 390

Example: Configuring TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

This example shows how to configure and verify Layer 3 TRM with PIM-SM for IPv4 multicast traffic when the RP is inside the BGP EVPN VXLAN fabric. The example uses the following topology:

Figure 43: TRM with PIM-SM when the RP is Inside the BGP EVPN VXLAN Fabric



The topology shows an EVPN VXLAN network, with two spine switches and three VTEPs, connected to an external network with three routers. VTEP 1 inside the BGP EVPN VXLAN fabric acts as the RP in this topology and Border VTEP connects the fabric to the external network through Router 1. The IPv4 multicast group is 226.1.1.1 in this topology. The following tables provide sample configurations for the devices in this topology:

Table 41: Configuring VTEP 1, Border VTEP, and VTEP 2 to Configure TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

VTEP 1	Border VTEP	VTEP 2
Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0	Border# **show running-config** hostname Border ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 ! vlan 2001 !	Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-famil ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !

VTEP 1	Border VTEP	VTEP 2
!	interface Loopback0	interface Loopback0
interface Loopback1	ip address 172.16.255.6	ip address 172.16.255.4
ip address 172.16.254.3	255.255.255.255	255.255.255.255
255.255.255.255	ip pim sparse-mode	ip pim sparse-mode
ip pim sparse-mode	ip ospf 1 area 0	ip ospf 1 area 0
ip ospf 1 area 0	!	!
!	interface Loopback1	interface Loopback1
interface Loopback255	ip address 172.16.254.6	ip address 172.16.254.4
vrf forwarding green	255.255.255.255	255.255.255.255
ip address 10.2.255.255	ip pim sparse-mode	ip pim sparse-mode
255.255.255.255	ip ospf 1 area 0	ip ospf 1 area 0
ip pim sparse-mode	!	!
!	interface GigabitEthernet1/0/1	interface GigabitEthernet1/0/1
interface GigabitEthernet1/0/1	no switchport	no switchport
no switchport	ip address 172.16.16.6	ip address 172.16.14.4
ip address 172.16.13.3	255.255.255.0	255.255.255.0
255.255.255.0	ip pim sparse-mode	ip pim sparse-mode
ip pim sparse-mode	ip ospf network point-to-point	ip ospf network point-to-point
ip ospf network point-to-point	ip ospf 1 area 0	ip ospf 1 area 0
ip ospf 1 area 0	!	!
!	interface GigabitEthernet1/0/2	interface GigabitEthernet1/0/2
interface GigabitEthernet1/0/2	no switchport	no switchport
no switchport	ip address 172.16.26.6	ip address 172.16.24.4
ip address 172.16.23.3	255.255.255.0	255.255.255.0
255.255.255.0	ip pim sparse-mode	ip pim sparse-mode
ip pim sparse-mode	ip ospf network point-to-point	ip ospf network point-to-point
ip ospf network point-to-point	ip ospf 1 area 0	ip ospf 1 area 0
ip ospf 1 area 0	!	!
!	interface GigabitEthernet1/0/5	interface
interface	switchport trunk allowed vlan	GigabitEthernet1/0/10
GigabitEthernet1/0/10	2001	switchport access vlan 102
switchport access vlan 101	switchport mode trunk	switchport mode access
switchport mode access	!	!
!	interface Vlan101	interface Vlan101
interface Vlan101	vrf forwarding green	vrf forwarding green
vrf forwarding green		ip address 10.1.101.1
ip address 10.1.101.1	ip pim sparse-mode	255.255.255.0
255.255.255.0	!	ip pim sparse-mode
ip pim sparse-mode	interface Vlan102	!
!	vrf forwarding green	interface Vlan102
interface Vlan102	ip address 10.1.102.1	vrf forwarding green
vrf forwarding green	255.255.255.0	ip address 10.1.102.1
ip address 10.1.102.1	ip pim sparse-mode	255.255.255.0
255.255.255.0	!	ip pim sparse-mode
ip pim sparse-mode	interface Vlan901	!
!	vrf forwarding green	interface Vlan901
interface Vlan901	ip unnumbered Loopback1	vrf forwarding green
vrf forwarding green	ip pim sparse-mode	ip unnumbered Loopback1
ip unnumbered Loopback1	no autostate	ip pim sparse-mode
ip pim sparse-mode	!	no autostate
no autostate		
!		

VTEP 1	Border VTEP	VTEP 2
interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute connected redistribute static exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family	interface Vlan2001 vrf forwarding green ip address 10.2.1.1 255.255.255.0 ip mtu 1500 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 2 vrf green redistribute bgp 65001 ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family	! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute connected redistribute static exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family !

VTEP 1	Border VTEP	VTEP 2
! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! end ! Leaf-01#	! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static redistribute ospf 2 match internal external 1 external 2 exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! end ! Border#	address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! end ! Leaf-02#

Table 42: Configuring Spine Switch 1 and Spine Switch 2 to Configure TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

Spine Switch 1	Spine Switch 2
Spine-01# **show running-config** hostname Spine-01 ! ip routing ! ip multicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.1 255.255.255.255 ip ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/4 no switchport ip address 172.16.16.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 neighbor 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 !	Spine-02# **show running-config** hostname Spine-02 ! ip routing ! ip multicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.2 255.255.255.255 ip ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/4 no switchport ip address 172.16.26.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp router-id 172.16.255.2 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 neighbor 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 !

Spine Switch 1	Spine Switch 2
address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-01#	address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-02#

Table 43: Configuring Router 1, Router 2, and Router 3 to Configure TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

Router 1	Router 2	Router 3
Router-01# **show running-config** hostname R1 ! ip multicast-routing distributed ! interface Loopback0 ip address 10.2.255.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1.2001 encapsulation dot1Q 2001 ip address 10.2.1.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 router ospf 1 router-id 10.2.255.1 ! ip pim rp-address 10.2.255.255 ! end ! R1#	Router-02# **show running-config** hostname R2 ! ip multicast-routing distributed ! interface Loopback0 ip address 10.2.255.2 255.255.255.255 ip pim sparse-mode ip igmp join-group 226.1.1.1 ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 10.2.255.2 ! ip pim rp-address 10.2.255.255 ! end ! R2#	Router-03# **show running-config** hostname R3 ! ip multicast-routing distributed ! interface Loopback0 ip address 10.2.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 10.2.255.3 ! ip pim rp-address 10.2.255.255 ! end ! R3#

Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

The following sections provide sample outputs for **show** commands to verify TRM with PIM-SM on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1
- Outputs to Verify the Configuration on VTEP 2
- Outputs to Verify the Configuration on Border VTEP
- Outputs to Verify the Configuration on Spine Switch 1

• Outputs to Verify the Configuration on Spine Switch 2

Outputs to Verify the Configuration on VTEP 1 (RP Inside the BGP EVPN VXLAN Fabric)

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Leaf-01# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 1d05h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 1d05h
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  1d05h
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  1d05h
Leaf-01#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
Leaf-01# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10102    172.16.254.4            7          10102    1d05h
nve1      10102    172.16.254.6            5          10102    1d05h
Leaf-01#
```

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on VTEP 1:

```
Leaf-01# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 111, main routing table version 111
7 network entries using 2128 bytes of memory
9 path entries using 1224 bytes of memory
5/5 BGP path/bestpath attribute entries using 1560 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
18 BGP extended community entries using 2396 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7492 total bytes of memory
BGP activity 140/45 prefixes, 240/112 paths, scan interval 60 secs
9 networks peaked at 12:22:24 Aug 6 2020 UTC (1d05h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    2104    1988      111    0    0 1d05h           2
172.16.255.2    4        65001    2099    1988      111    0    0 1d05h           2
Leaf-01#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **rp mapping** command on VTEP 1:

```
Leaf-01# show ip pim vrf green rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
```

```
Leaf-01#
```

The following example shows the output for the **show ip routing vrf** command on VTEP 1:

```
Leaf-01# show ip routing vrf green 10.2.255.255
Routing Table: green
Routing entry for 10.2.255.255/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Redistributing via bgp 65001
  Advertised by bgp 65001
  Routing Descriptor Blocks:
  * directly connected, via Loopback255
      Route metric is 0, traffic share count is 1
Leaf-01#
```

The following example shows the output for the **show ip igmp vrf** *vrf-name* **groups** command on VTEP 1:

```
Leaf-01# show ip igmp vrf green groups
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter   Group Accounted
226.1.1.1        Vlan102                  1d05h     00:02:50  10.1.102.12
224.0.1.40       Vlan901                  1d05h     00:02:03  172.16.254.4
Leaf-01#
```

The following example shows the output for the **show ip mroute vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 1d01h/stopped, RP 10.2.255.255, flags: SJCGx
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan101, Forward/Sparse, 1d01h/00:02:33
    Vlan901, Forward/Sparse, 03:54:15/stopped

(10.2.255.1, 226.1.1.1), 00:01:13/00:01:50, flags: Tgx
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list:
    Vlan101, Forward/Sparse, 00:01:13/00:02:33

(10.1.102.12, 226.1.1.1), 00:01:36/00:01:24, flags: Tgx
```

```
  Incoming interface: Vlan901, RPF nbr 172.16.254.4
  Outgoing interface list:
    Vlan101, Forward/Sparse, 00:01:36/00:02:33

(*, 224.0.1.40), 1d05h/00:02:09, RP 10.2.255.255, flags: SJCLGx
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback901, Forward/Sparse, 1d05h/00:02:09
    Vlan901, Forward/Sparse, 03:54:15/stopped
Leaf-01#
```

The following example shows the output for the **show ip mfib vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ip mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 2/2/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel6 Flags: A
   Loopback901 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (*,226.1.1.1) Flags: C HW
   SW Forwarding: 1/0/100/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel6 Flags: A
   Vlan101 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/1    Rate: 0 pps
 (10.1.102.12,226.1.1.1) Flags: HW
   SW Forwarding: 2/0/100/0, Other: 0/0/0
   HW Forwarding:   44/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan101 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (10.2.255.1,226.1.1.1) Flags: HW
   SW Forwarding: 5/0/100/0, Other: 12576/1/12575
   HW Forwarding:   3801/1/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan901 Flags: SP
   Vlan101 Flags: F NS
```

```
     Pkts: 0/0/5    Rate: 0 pps
Leaf-01#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on VTEP 1:

```
Leaf-01# show bgp ipv4 mvpn all
BGP table version is 94, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 *>i  [5][1:1][10.2.255.1][226.1.1.1]/18
                      172.16.255.6             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 *>   [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                      0.0.0.0                            32768 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                      0.0.0.0                            32768 ?
Leaf-01#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Leaf-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:10
```

```
(172.16.254.6, 239.1.1.1), 00:01:11/00:01:48, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:01:11/00:01:48

(172.16.254.3, 239.1.1.1), 00:01:37/00:01:22, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0, Registering
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:37/00:02:51, A

(172.16.254.4, 239.1.1.1), 04:17:32/00:02:31, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 04:17:32/00:00:27

(*, 224.0.1.40), 1d05h/00:02:12, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d05h/00:02:12

(*, 225.0.0.102), 1d05h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:10

(172.16.254.4, 225.0.0.102), 1d05h/00:01:20, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:32

(172.16.254.6, 225.0.0.102), 1d05h/00:02:44, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:10

(*, 225.0.0.101), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:10

(172.16.254.3, 225.0.0.101), 1d05h/00:02:36, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:03:20
Leaf-01#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Leaf-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup
```

```
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/114/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 13/0/127/0, Other: 2/2/0
   HW Forwarding:   12525/0/165/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   9155/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   3762/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 2/2/0
   HW Forwarding:   15/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 2/1/1
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
   Tunnel4 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   7707/0/167/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
```

```
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 2/0/150/0, Other: 0/0/0
   HW Forwarding:   68/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
Leaf-01#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 431

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Leaf-02# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 1d05h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 1d05h
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  1d05h
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  1d05h
Leaf-02#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
Leaf-02# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            6          10101    1d05h
nve1      10102    172.16.254.6            5          10102    1d05h
Leaf-02#
```

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on VTEP 2:

```
Leaf-02# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 62, main routing table version 62
7 network entries using 2128 bytes of memory
9 path entries using 1224 bytes of memory
4/4 BGP path/bestpath attribute entries using 1248 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2372 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7156 total bytes of memory
BGP activity 121/28 prefixes, 202/77 paths, scan interval 60 secs
9 networks peaked at 12:22:24 Aug 6 2020 UTC (16:43:21.423 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1229    1151       62    0    0 16:57:50        2
172.16.255.2    4        65001    1227    1152       62    0    0 16:57:51        2
```

```
Leaf-02#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **rp mapping** command on VTEP 2:

```
Leaf-02# show ip pim vrf green rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
Leaf-02#
```

The following example shows the output for the **show ip routing vrf** command on VTEP 2:

```
Leaf-02# show ip routing vrf green 10.2.255.255
Routing Table: green
Routing entry for 10.2.255.255/32
  Known via "bgp 65001", distance 200, metric 0, type internal
  Last update from 172.16.254.3 on Vlan901, 03:59:59 ago
  Routing Descriptor Blocks:
  * 172.16.254.3 (default), from 172.16.255.1, 03:59:59 ago, via Vlan901
      opaque_ptr 0x7F65B8B9E9F0
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none
Leaf-02#
```

The following example shows the output for the **show ip igmp vrf** *vrf-name* **groups** command on VTEP 2:

```
Leaf-02# show ip igmp vrf green groups
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter   Group Accounted
226.1.1.1        Vlan102                  16:58:00  00:02:11  10.1.102.12
224.0.1.40       Vlan901                  16:58:37  00:02:33  172.16.254.4
Leaf-02#
```

The following example shows the output for the **show ip mroute vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(*, 226.1.1.1), 1d05h/stopped, RP 10.2.255.255, flags: SJCFg
  Incoming interface: Vlan901, RPF nbr 172.16.254.3
  Outgoing interface list:
    Vlan102, Forward/Sparse, 1d05h/00:02:50

(10.2.255.1, 226.1.1.1), 00:06:57/00:02:09, flags: TgQ
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list:
    Vlan102, Forward/Sparse, 00:06:57/00:02:50

(10.1.102.12, 226.1.1.1), 00:07:21/00:01:45, flags: FTGqx
  Incoming interface: Vlan102, RPF nbr 0.0.0.0
  Outgoing interface list:
    Vlan901, Forward/Sparse, 00:07:21/stopped

(*, 224.0.1.40), 1d05h/00:02:03, RP 10.2.255.255, flags: SJPCLgx
  Incoming interface: Vlan901, RPF nbr 172.16.254.3
  Outgoing interface list: Null
Leaf-02#
```

The following example shows the output for the **show ip mfib vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ip mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A IC NS
 (*,226.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   3/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan102 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (10.1.102.12,226.1.1.1) Flags: HW
   SW Forwarding: 2/0/100/0, Other: 0/0/0
   HW Forwarding:   215/0/118/0, Other: 0/0/0
   Vlan102 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/2    Rate: 0 pps
 (10.2.255.1,226.1.1.1) Flags: HW
   SW Forwarding: 2/0/100/0, Other: 0/0/0
Leaf-02#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on VTEP 2:

```
Leaf-02# show bgp ipv4 mvpn all
BGP table version is 94, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 *>i  [5][1:1][10.2.255.1][226.1.1.1]/18
                     172.16.255.6             0    100      0 ?
 * i                  172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 *>   [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                     0.0.0.0                            32768 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                     0.0.0.0                            32768 ?
Leaf-02#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Leaf-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:59

(172.16.254.6, 239.1.1.1), 00:06:55/00:01:59, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:06:55/00:02:04
```

```
(172.16.254.4, 239.1.1.1), 04:23:16/00:03:29, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 04:23:16/00:02:57

(*, 224.0.1.40), 1d05h/00:02:02, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d05h/00:02:02
Leaf-02#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Leaf-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/170/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   12469/0/177/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/224/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 2/0/163/0, Other: 3/1/2
   HW Forwarding:   9233/0/164/0, Other: 0/0/0
```

```
    Null0 Flags: A
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 1/0/206/0, Other: 0/0/0
    HW Forwarding:   3767/0/163/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   8/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 1/0/150/0, Other: 22/18/4
    HW Forwarding:   7870/0/156/0, Other: 0/0/0
    Null0 Flags: A
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 2/0/150/0, Other: 0/0/0
    HW Forwarding:   412/1/168/1, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/2    Rate: 0 pps
Leaf-02#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 431

Outputs to Verify the Configuration on Border VTEP

The following example shows the output for the **show nve peers** command on Border VTEP:

```
Border# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 1d05h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 1d05h
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  1d05h
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  1d05h
Border#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on Border VTEP:

```
Border# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            6          10101    1d05h
nve1      10102    172.16.254.4            7          10102    1d05h
Border#
```

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on Border VTEP:

```
Border# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.6, local AS number 65001
BGP table version is 102, main routing table version 102
6 network entries using 1824 bytes of memory
8 path entries using 1088 bytes of memory
5/5 BGP path/bestpath attribute entries using 1560 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
20 BGP extended community entries using 2706 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7362 total bytes of memory
BGP activity 133/39 prefixes, 265/144 paths, scan interval 60 secs
8 networks peaked at 12:14:22 Aug 6 2020 UTC (1d05h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    2114    1995      102    0    0 1d05h           2
172.16.255.2    4        65001    2112    1990      102    0    0 1d05h           2
Border#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **rp mapping** command on Border VTEP:

```
Border# show ip pim vrf green rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
Border#
```

The following example shows the output for the **show ip routing vrf** *vrf-name* command on Border VTEP:

```
Border# show ip routing vrf green 10.2.255.255
Routing Table: green
Routing entry for 10.2.255.255/32
  Known via "bgp 65001", distance 200, metric 0, type internal
  Redistributing via ospf 2
  Advertised by ospf 2 subnets
  Last update from 172.16.254.3 on Vlan901, 04:02:51 ago
  Routing Descriptor Blocks:
  * 172.16.254.3 (default), from 172.16.255.1, 04:02:51 ago, via Vlan901
      opaque_ptr 0x7FEF6836D190
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none
Border#
```

The following example shows the output for the **show ip igmp vrf** *vrf-name* **groups** command on Border VTEP:

```
Border# show ip igmp vrf green groups
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter   Group Accounted
224.0.1.40       Vlan901                  1d05h     00:01:58  172.16.254.6
Border#
```

The following example shows the output for the **show ip mroute vrf** *vrf-name* command on Border VTEP:

```
Border# show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 1d05h/00:03:07, RP 10.2.255.255, flags: SJgx
  Incoming interface: Vlan901, RPF nbr 172.16.254.3
  Outgoing interface list:
    Vlan2001, Forward/Sparse, 04:02:51/00:03:07

(10.2.255.1, 226.1.1.1), 00:09:51/00:01:02, flags: TGqx
  Incoming interface: Vlan2001, RPF nbr 10.2.1.2
  Outgoing interface list:
    Vlan901, Forward/Sparse, 00:09:49/stopped

(10.1.102.12, 226.1.1.1), 00:10:12/00:03:09, flags: TgQx
  Incoming interface: Vlan901, RPF nbr 172.16.254.4
  Outgoing interface list:
    Vlan2001, Forward/Sparse, 00:10:12/00:03:07

(*, 224.0.1.40), 1d05h/00:03:10, RP 10.2.255.255, flags: SJCLgx
  Incoming interface: Vlan901, RPF nbr 172.16.254.3
  Outgoing interface list:
    Vlan2001, Forward/Sparse, 04:02:51/00:03:10
Border#
```

The following example shows the output for the **show ip mfib vrf** *vrf-name* command on Border VTEP:

```
Border# show ip mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
```

```
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A IC NS
   Vlan2001 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,226.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   6/0/122/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan2001 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (10.1.102.12,226.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   304/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan2001 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (10.2.255.1,226.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 1/0/1
   HW Forwarding:   582/1/122/0, Other: 0/0/0
   Vlan2001 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
Border#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on Border VTEP:

```
Border# show bgp ipv4 mvpn allBGP table version is 102, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 *>   [5][1:1][10.2.255.1][226.1.1.1]/18
                      0.0.0.0                            32768 ?
 *>   [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                      0.0.0.0                            32768 ?
 *>   [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                      0.0.0.0                            32768 ?
 *>i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                      172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                      0.0.0.0                            32768 ?
Border#
```

The following example shows the output for the **show ip mroute** command on Border VTEP:

```
Border# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:56

(172.16.254.6, 239.1.1.1), 00:09:47/00:02:24, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:09:47/00:02:33

(172.16.254.4, 239.1.1.1), 04:26:08/00:02:10, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 04:26:08/00:00:51

(*, 224.0.1.40), 1d05h/00:02:56, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d05h/00:02:56

(*, 225.0.0.102), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:56

(172.16.254.4, 225.0.0.102), 1d05h/00:01:27, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:56

(172.16.254.6, 225.0.0.102), 1d05h/00:01:53, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:43, A

(*, 225.0.0.101), 1d05h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:56

(172.16.254.3, 225.0.0.101), 1d05h/00:01:10, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
```

```
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:08
Border#
```

The following example shows the output for the **show ip mfib** command on Border VTEP:

```
Border# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   8/0/146/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 7/0/125/0, Other: 0/0/0
   HW Forwarding:   12570/0/177/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/7    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   9199/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 17/0/174/0, Other: 10/9/1
   HW Forwarding:   3789/0/151/0, Other: 0/0/0
   Null0 Flags: A
```

```
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/16    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   9/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   7961/0/167/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 2/0/150/0, Other: 2/2/0
   HW Forwarding:   580/1/156/1, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
Border#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 431

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on Spine Switch 1:

```
Spine-01# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 204, main routing table version 204
6 network entries using 1824 bytes of memory
16 path entries using 2176 bytes of memory
4/4 BGP path/bestpath attribute entries using 1216 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2356 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7716 total bytes of memory
BGP activity 266/218 prefixes, 3029/2926 paths, scan interval 60 secs
8 networks peaked at 12:20:11 Aug 6 2020 UTC (1d05h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    1985    2115      204    0    0 1d05h           2
172.16.255.4    4        65001    1995    2111      204    0    0 1d05h           4
172.16.255.6    4        65001    1999    2118      204    0    0 1d05h           4
Spine-01#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 1:

```
Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings
```

```
Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-01#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on Spine Switch 1:

```
Spine-01# show bgp ipv4 mvpn all
BGP table version is 204, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 * i  [5][1:1][10.2.255.1][226.1.1.1]/18
                      172.16.255.6             0    100      0 ?
 *>i                   172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
 * i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 * i  [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                      172.16.255.6             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Spine-01#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 1:

```
Spine-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
```

```
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 04:29:40/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 239.1.1.1), 00:13:17/00:02:24, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(172.16.254.4, 239.1.1.1), 04:27:38/00:02:41, flags: PTA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.14.4
  Outgoing interface list: Null

(*, 224.0.1.40), 1w0d/00:02:43, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1w0d/00:02:43

(*, 225.0.0.102), 1w0d/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null
Spine-01#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 1:

```
Spine-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 83/83/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 1/0/206/0, Other: 282/0/282
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: NS
```

```
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/4 Flags: NS
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 4/3/1
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   7/0/158/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/4 Flags: NS
Spine-01#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 431

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on Spine Switch 2:

```
Spine-02# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 164, main routing table version 164
6 network entries using 1824 bytes of memory
16 path entries using 2176 bytes of memory
4/4 BGP path/bestpath attribute entries using 1216 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2356 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7716 total bytes of memory
BGP activity 297/249 prefixes, 3131/3028 paths, scan interval 60 secs
8 networks peaked at 12:20:59 Aug 6 2020 UTC (1d05h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    1988    2114      164    0    0 1d05h           2
172.16.255.4    4        65001    1998    2110      164    0    0 1d05h           4
172.16.255.6    4        65001    1996    2119      164    0    0 1d05h           4
Spine-02#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 2:

```
Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
```

```
    RP: 172.16.255.255 (?)
Spine-02#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on Spine Switch 2:

```
Spine-02# show bgp ipv4 mvpn all
BGP table version is 164, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 * i  [5][1:1][10.2.255.1][226.1.1.1]/18
                      172.16.255.6             0    100      0 ?
 *>i                   172.16.255.6             0    100      0 ?
 *>i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
 *>i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
 * i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 * i  [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                      172.16.255.6             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Spine-02#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 2:

```
Spine-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode
```

```
(*, 239.1.1.1), 3d12h/00:03:14, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:31
    GigabitEthernet1/0/1, Forward/Sparse, 1d05h/00:03:14
    GigabitEthernet1/0/4, Forward/Sparse, 1d05h/00:03:09

(172.16.254.6, 239.1.1.1), 00:15:48/00:01:26, flags: T
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:15:48/00:03:24
    GigabitEthernet1/0/2, Forward/Sparse, 00:15:48/00:03:26

(172.16.254.4, 239.1.1.1), 04:32:09/00:01:28, flags: T
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 04:32:09/00:03:09
    GigabitEthernet1/0/1, Forward/Sparse, 04:32:09/00:03:14

(*, 224.0.1.40), 1w0d/00:03:29, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:48
    GigabitEthernet1/0/1, Forward/Sparse, 1d05h/00:03:29
    GigabitEthernet1/0/4, Forward/Sparse, 1d05h/00:03:17
    Loopback2, Forward/Sparse, 1w0d/00:02:34

(*, 225.0.0.102), 1w0d/00:03:28, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:55
    GigabitEthernet1/0/1, Forward/Sparse, 1d05h/00:03:00
    GigabitEthernet1/0/4, Forward/Sparse, 1d05h/00:03:28

(172.16.254.4, 225.0.0.102), 1d05h/00:02:09, flags: MT
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d05h/00:03:28
    GigabitEthernet1/0/1, Forward/Sparse, 1d05h/00:03:17

(172.16.254.6, 225.0.0.102), 1d05h/00:01:40, flags: MT
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:59
    GigabitEthernet1/0/1, Forward/Sparse, 1d05h/00:03:05

(*, 225.0.0.101), 3d12h/00:03:21, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:33
    GigabitEthernet1/0/1, Forward/Sparse, 1d05h/00:03:21
    GigabitEthernet1/0/4, Forward/Sparse, 1d05h/00:02:47

(172.16.254.3, 225.0.0.101), 1d05h/00:02:05, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:57
    GigabitEthernet1/0/4, Forward/Sparse, 1d05h/00:02:47
Spine-02#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 2:

```
Spine-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 9/0/112/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 4/0/132/0, Other: 0/0/0
   HW Forwarding:   12607/0/177/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 27/0/101/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   9232/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
```

```
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   3789/0/163/0, Other: 0/0/0
    GigabitEthernet1/0/4 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 10/0/150/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/4    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/4    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/4    Rate: 0 pps
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   8144/0/167/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   942/1/168/1, Other: 0/0/0
    GigabitEthernet1/0/4 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
Spine-02#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 431

Example: Configuring TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

This example shows how to configure and verify Layer 3 TRM with PIM-SM for IPv4 and IPv6 multicast traffic when the RP is inside the BGP EVPN VXLAN fabric. The example uses the topology in the TRM with PIM-SM when the RP is inside the BGP EVPN VXLAN fabric figure.

The topology shows an EVPN VXLAN network, with two spine switches and three VTEPs, connected to an external network with three routers. VTEP 1 inside the BGP EVPN VXLAN fabric acts as the RP in this topology and Border VTEP connects the fabric to the external network through Router 1. The IPv4 multicast group is 226.1.1.1 and the IPv6 multicast group is FF06:1::1 in this topology. The following tables provide sample configurations for the devices in this topology:

Table 44: Configuring VTEP 1, Border VTEP, and VTEP 2 to Configure TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

VTEP 1	Border VTEP	VTEP 2
Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !	Border# **show running-config** hostname Border ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !	Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !

VTEP 1	Border VTEP	VTEP 2
interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback255 vrf forwarding green ip address 10.2.255.255 255.255.255.255 ip pim sparse-mode ipv6 address FC00:2:255::255/128 ipv6 enable ! interface Loopback901 vrf forwarding green ip address 10.1.255.1 255.255.255.255 ip pim sparse-mode ipv6 address FC00:1:255::1/128 ipv6 enable ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:101::1/64 ipv6 enable	vlan 2001 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback901 vrf forwarding green ip address 10.1.255.4 255.255.255.255 ip pim sparse-mode ipv6 address FC00:1:255::4/128 ipv6 enable ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 switchport trunk allowed vlan 2001 switchport mode trunk ! interface GigabitEthernet1/0/10 switchport access vlan 102 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:101::1/64 ipv6 enable !	interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback901 vrf forwarding green ip address 10.1.255.2 255.255.255.255 ip pim sparse-mode ipv6 address FC00:1:255::2/128 ipv6 enable ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ipv6 address FC00:1:102::1/64 ipv6 enable

VTEP 1	Border VTEP	VTEP 2
!	interface Vlan102	
interface Vlan102	vrf forwarding green	!
vrf forwarding green	ip address 10.1.102.1 255.255.255.0	interface Vlan901
ip address 10.1.102.1 255.255.255.0	ip pim sparse-mode	vrf forwarding green
ip pim sparse-mode	ipv6 address FC00:1:102::1/64	ip unnumbered Loopback1
ipv6 address FC00:1:102::1/64	ipv6 enable	ip pim sparse-mode
ipv6 enable	!	ipv6 enable
!	interface Vlan901	no autostate
interface Vlan901	vrf forwarding green	!
vrf forwarding green	ip unnumbered Loopback1	interface nve1
ip unnumbered Loopback1	ip pim sparse-mode	no ip address
ip pim sparse-mode	ipv6 enable	source-interface Loopback1
ipv6 enable	no autostate	host-reachability protocol bgp
no autostate	!	member vni 10101 mcast-group
!	interface Vlan2001	225.0.0.101
interface nve1	vrf forwarding green	member vni 50901 vrf green
no ip address	ip address 10.2.1.1 255.255.255.0	member vni 10102 mcast-group
source-interface Loopback1	ip mtu 1500	225.0.0.102
host-reachability protocol bgp	ip pim sparse-mode	!
member vni 10101 mcast-group	ip ospf network point-to-point	router ospf 1
225.0.0.101	ip ospf 2 area 0	router-id 172.16.255.4
member vni 50901 vrf green	ipv6 address FC00:2:1::1/64	!
member vni 10102 mcast-group	ipv6 enable	router bgp 65001
225.0.0.102	ipv6 mtu 1500	bgp log-neighbor-changes
!	ospfv3 network point-to-point	no bgp default ipv4-unicast
router ospf 1	ospfv3 1 ipv6 area 0	neighbor 172.16.255.1 remote-as 65001
router-id 172.16.255.3	!	neighbor 172.16.255.1 update-source
!	interface nve1	Loopback0
router bgp 65001	no ip address	neighbor 172.16.255.2 remote-as 65001
bgp log-neighbor-changes	source-interface Loopback1	neighbor 172.16.255.2 update-source
no bgp default ipv4-unicast	host-reachability protocol bgp	Loopback0
neighbor 172.16.255.1 remote-as 65001	member vni 10101 mcast-group	!
neighbor 172.16.255.1 update-source	225.0.0.101	address-family ipv4
Loopback0	member vni 50901 vrf green	redistribute connected
neighbor 172.16.255.2 remote-as 65001	member vni 10102 mcast-group	redistribute static
neighbor 172.16.255.2 update-source	225.0.0.102	exit-address-family
Loopback0	!	!
!	router ospfv3 1	address-family ipv4 mvpn
address-family ipv4	!	neighbor 172.16.255.1 activate
redistribute connected	address-family ipv6 unicast vrf green	neighbor 172.16.255.1 send-community
redistribute static	redistribute bgp 65001	both
exit-address-family	exit-address-family	neighbor 172.16.255.2 activate
!	!	neighbor 172.16.255.2 send-community
address-family ipv4 mvpn	router ospf 2 vrf green	both
neighbor 172.16.255.1 activate	redistribute bgp 65001	exit-address-family
neighbor 172.16.255.1 send-community	!	!
both	router ospf 1	address-family ipv6 mvpn
neighbor 172.16.255.2 activate	router-id 172.16.255.6	neighbor 172.16.255.1 activate
neighbor 172.16.255.2 send-community	!	neighbor 172.16.255.1 send-community
both	router bgp 65001	both
exit-address-family	bgp log-neighbor-changes	neighbor 172.16.255.2 activate
!	no bgp default ipv4-unicast	neighbor 172.16.255.2 send-community
	neighbor 172.16.255.1 remote-as 65001	both
	neighbor 172.16.255.1 update-source	exit-address-family
	Loopback0	
	neighbor 172.16.255.2 remote-as 65001	
	neighbor 172.16.255.2 update-source	
	Loopback0	

VTEP 1	Border VTEP	VTEP 2
address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end ! Leaf-01#	! address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static redistribute ospf 2 match internal external 1 external 2 exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute ospf 1 include-connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ! end ! Border#	! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end ! Leaf-02#

Table 45: Configuring Spine Switch 1 and Spine Switch 2 to Configure TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

Spine Switch 1

```
Spine-01# show running-config
hostname Spine-01
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.14.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/4
no switchport
ip address 172.16.16.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.1
!
router bgp 65001
bgp router-id 172.16.255.1
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0
!
```

Spine Switch 2

```
Spine-02# show running-config
hostname Spine-02
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.23.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.24.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/4
no switchport
ip address 172.16.26.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.2
!
router bgp 65001
bgp router-id 172.16.255.2
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0
!
```

Spine Switch 1	Spine Switch 2
address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-01#	address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-02#

Table 46: Configuring Router 1, Router 2, and Router 3 to Configure TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

Router 1	Router 2	Router 3
R1# **show running-config** hostname R1 ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! interface Loopback0 ip address 10.2.255.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ipv6 address FC00:2:255::1/128 ipv6 enable ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:12::1/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:13::1/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface GigabitEthernet0/0/1.2001 encapsulation dot1Q 2001 ip address 10.2.1.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ipv6 address FC00:2:1::2/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 !	R2# **show running-config** hostname R2 ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! interface Loopback0 ip address 10.2.255.2 255.255.255.255 ip pim sparse-mode ip igmp join-group 226.1.1.1 ip ospf 1 area 0 ipv6 address FC00:2:255::2/128 ipv6 enable ipv6 mld join-group FF06:1::1 ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:12::2/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:23::2/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! router ospfv3 1 ! address-family ipv6 unicast exit-address-family ! router ospf 1 router-id 10.2.255.2 ! ip pim rp-address 10.2.255.255	R3# **show running-config** hostname R3 ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! interface Loopback0 ip address 10.2.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ipv6 address FC00:2:255::3/128 ipv6 enable ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:13::3/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:23::3/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! router ospfv3 1 ! address-family ipv6 unicast exit-address-family ! router ospf 1 router-id 10.2.255.3 ! ip pim rp-address 10.2.255.255 !

Router 1	Router 2	Router 3
router ospfv3 1 ! address-family ipv6 unicast exit-address-family ! router ospf 1 router-id 10.2.255.1 ! ip pim rp-address 10.2.255.255 ! ipv6 pim rp-address FC00:2:255::255 ! end ! R1#	! ipv6 pim rp-address FC00:2:255::255 ! end ! R2#	ipv6 pim rp-address FC00:2:255::255 ! end ! R3#

Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric

The following sections provide sample outputs for **show** commands to verify TRM with PIM-SM on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1
- Outputs to Verify the Configuration on VTEP 2
- Outputs to Verify the Configuration on Border VTEP
- Outputs to Verify the Configuration on Spine Switch 1
- Outputs to Verify the Configuration on Spine Switch 2

Outputs to Verify the Configuration on VTEP 1 (RP Inside the BGP EVPN VXLAN Fabric)

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Leaf-01# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 1d05h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 1d05h
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  1d05h
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  1d05h
Leaf-01#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
Leaf-01# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10102    172.16.254.4            7          10102    1d05h
nve1      10102    172.16.254.6            5          10102    1d05h
Leaf-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on VTEP 1:

```
Leaf-01# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 60, main routing table version 60
5 network entries using 1960 bytes of memory
8 path entries using 1280 bytes of memory
4/4 BGP path/bestpath attribute entries using 1248 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
18 BGP extended community entries using 2396 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7068 total bytes of memory
BGP activity 139/41 prefixes, 275/138 paths, scan interval 60 secs
5 networks peaked at 15:46:09 Aug 6 2020 UTC (1d02h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    2135    2003       60    0    0 1d05h           3
172.16.255.2    4        65001    2131    2003       60    0    0 1d05h           3
Leaf-01#
```

The following example shows the output for the **show ipv6 pim vrf** *vrf-name* **group-map** command on VTEP 1:

```
Leaf-01# show ipv6 pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Tu7,FC00:2:255::255 (us)
    Info source: Static
    Uptime: 1d02h, Groups: 1
Leaf-01#
```

The following example shows the output for the **show ipv6 routing vrf** command on VTEP 1:

```
Leaf-01# show ipv6 routing vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "connected", distance 0, metric 0, type receive, connected
  Redistributing via bgp 65001
  Route count is 1/1, share count 0
  Routing paths:
    receive via Loopback255
      Last updated 04:21:51 ago
Leaf-01#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 1:

```
Leaf-01# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan101
   1d02h     00:02:28
Leaf-01#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 1d02h/00:03:13, RP FC00:2:255::255, flags: SCJG
  Incoming interface: Tunnel7
  RPF nbr: FC00:2:255::255
  Immediate Outgoing interface list:
    Vlan101, Forward, 1d02h/00:03:13
    Vlan901, Forward, 04:21:51/never

(FC00:1:102::12, FF06:1::1), 00:02:17/00:03:04, flags: SJTg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.4
  Immediate Outgoing interface list:
    Vlan101, Forward, 00:02:13/00:03:23

(FC00:2:255::1, FF06:1::1), 00:01:24/00:03:04, flags: SJTg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.6
  Immediate Outgoing interface list:
    Vlan101, Forward, 00:01:20/00:03:13
Leaf-01#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF00::/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 412/412/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel7 Flags: NS
```

```
(*,FF00::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF02::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF06:1::1) Flags: C HW
  SW Forwarding: 2/0/100/0, Other: 0/0/0
  HW Forwarding:   4/0/126/0, Other: 0/0/0
  Tunnel7 Flags: A NS
  Vlan101 Flags: F NS
    Pkts: 0/0/2    Rate: 0 pps
  Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
    Pkts: 0/0/2    Rate: 0 pps
(FC00:1:102::12,FF06:1::1) Flags: HW
  SW Forwarding: 1/0/100/0, Other: 0/0/0
  HW Forwarding:   64/0/126/0, Other: 0/0/0
  Vlan901, VXLAN Decap Flags: A
  Vlan101 Flags: F NS
    Pkts: 0/0/1    Rate: 0 pps
(FC00:2:255::1,FF06:1::1) Flags: HW
  SW Forwarding: 2/0/100/0, Other: 0/0/0
  HW Forwarding:   38/0/126/0, Other: 0/0/0
  Vlan901, VXLAN Decap Flags: A
  Vlan101 Flags: F NS
    Pkts: 0/0/2    Rate: 0 pps
(*,FF10::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF12::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF20::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF22::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF30::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF32::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF33::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF34::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF35::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF36::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF37::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF38::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF39::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3A::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
```

```
(*,FF3B::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3C::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3D::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3E::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3F::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF40::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF42::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF50::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF52::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF60::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF62::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF70::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF72::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF80::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF82::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF90::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF92::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFA0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFA2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFB0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFB2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFC0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFC2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFD0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFD2::/16) Flags:
```

```
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFE0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFE2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFF0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFF2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
Leaf-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 1:

```
Leaf-01# show bgp ipv6 mvpn all
BGP table version is 60, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
                     172.16.255.6             0    100      0 ?
 *>i                  172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 *>   [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                     ::                                    32768 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                     ::                                    32768 ?
Leaf-01#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Leaf-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
```

```
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:34

(172.16.254.3, 239.1.1.1), 00:02:17/00:02:05, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0, Registering
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:02:17/00:03:09, A

(172.16.254.6, 239.1.1.1), 00:28:47/00:02:22, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:28:47/00:01:12

(172.16.254.4, 239.1.1.1), 04:45:08/00:01:03, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 04:45:08/00:02:51

(*, 224.0.1.40), 1d05h/00:02:38, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d05h/00:02:38

(*, 225.0.0.102), 1d05h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:34

(172.16.254.4, 225.0.0.102), 1d05h/00:02:33, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:00:56

(172.16.254.6, 225.0.0.102), 1d05h/00:01:12, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:34

(*, 225.0.0.101), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:01:34

(172.16.254.3, 225.0.0.101), 1d05h/00:03:17, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:03:17
Leaf-01#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Leaf-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
```

```
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/114/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 13/0/127/0, Other: 2/2/0
   HW Forwarding:   12686/0/165/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   9299/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   3817/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 4/4/0
   HW Forwarding:   15/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 6/5/1
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Null0 Flags: A NS
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
```

```
   Tunnel4 Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   8525/0/167/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 2/0/150/0, Other: 0/0/0
   HW Forwarding:   1629/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
Leaf-01#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 464

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Leaf-02# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 1d05h
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 1d05h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 1d05h
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  1d05h
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  1d05h
Leaf-02#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
Leaf-02# show l2vpn evpn peers vxlan

Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            6          10101    1d05h
nve1      10102    172.16.254.6            5          10102    1d05h
Leaf-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on VTEP 2:

```
Leaf-02# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 85, main routing table version 85
6 network entries using 2352 bytes of memory
8 path entries using 1280 bytes of memory
5/5 BGP path/bestpath attribute entries using 1560 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
18 BGP extended community entries using 2396 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7772 total bytes of memory
```

```
BGP activity 145/47 prefixes, 249/117 paths, scan interval 60 secs
6 networks peaked at 15:50:41 Aug 6 2020 UTC (1d02h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    2143    2019       85    0    0 1d05h           2
172.16.255.2    4        65001    2139    2019       85    0    0 1d05h           2
Leaf-02#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **group-map** command on VTEP 2:

```
Leaf-02# show ip pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Vl901,::FFFF:172.16.254.3
    Info source: Static
    Uptime: 1d05h, Groups: 1
Leaf-02#
```

The following example shows the output for the **show ip routing vrf** command on VTEP 2:

```
Leaf-02# show ip routing vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "bgp 65001", distance 200, metric 0, type internal
  Route count is 1/1, share count 0
  Routing paths:
    172.16.254.3%default, Vlan901%default
      From AC10:FF01::
      opaque_ptr 0x7F65BA333EC0
      Last updated 04:26:58 ago
Leaf-02#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 2:

```
Leaf-02# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan102
   1d05h     00:03:53
Leaf-02#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
```

```
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 1d05h/never, RP FC00:2:255::255, flags: SCJg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.3
  Immediate Outgoing interface list:
    Vlan102, Forward, 1d05h/never

(FC00:1:102::12, FF06:1::1), 00:07:24/00:03:28, flags: SFJTGq
  Incoming interface: Vlan102
  RPF nbr: FE80::46D3:CAFF:FE28:6CC5
  Immediate Outgoing interface list:
    Vlan901, Forward, 00:07:24/never

(FC00:2:255::1, FF06:1::1), 00:06:31/00:00:32, flags: SJTgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.6
  Inherited Outgoing interface list:
    Vlan102, Forward, 1d05h/never
Leaf-02#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF00::/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF00::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF02::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 10/10/0
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   3/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan102 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:102::12,FF06:1::1) Flags: HW
   SW Forwarding: 2/0/100/0, Other: 0/0/0
   HW Forwarding:   217/0/118/0, Other: 0/0/0
   Vlan102 Flags: A F
     Pkts: 0/0/0    Rate: 0 pps
```

```
  Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
    Pkts: 0/0/1    Rate: 0 pps
(FC00:2:255::1,FF06:1::1) Flags: HW
  SW Forwarding: 2/0/100/0, Other: 0/0/0
  HW Forwarding:   191/0/126/0, Other: 0/0/0
  Vlan901, VXLAN Decap Flags: A
  Vlan102 Flags: F NS
    Pkts: 0/0/2    Rate: 0 pps
(*,FF10::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF12::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF20::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF22::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF30::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF32::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF33::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF34::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF35::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF36::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF37::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF38::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF39::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3A::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3B::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3C::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3D::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3E::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3F::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF40::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF42::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF50::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF52::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF60::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF62::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF70::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF72::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF80::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF82::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF90::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF92::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFA0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFA2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFB0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFB2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFC0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFC2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFD0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFD2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFE0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFE2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFF0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFF2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
Leaf-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 2:

```
Leaf-02# show bgp ipv6 mvpn all
BGP table version is 85, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>   [5][1:1][FC00:1:102::12][FF06:1::1]/42
                       ::                                    32768 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
                       172.16.255.6             0    100      0 ?
 *>i                    172.16.255.6             0    100      0 ?
 *>   [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                       ::                                    32768 ?
 *>i  [7][1:1][65001][FC00:1:102::12][FF06:1::1]/46
                       172.16.255.3             0    100      0 ?
 *>   [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                       ::                                    32768 ?
Route Distinguisher: 172.16.254.4:102
 * i  [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                       172.16.255.3             0    100      0 ?
 *>i                    172.16.255.3             0    100      0 ?
Leaf-02#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Leaf-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:00

(172.16.254.6, 239.1.1.1), 00:33:54/00:01:36, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:33:54/00:02:05

(172.16.254.4, 239.1.1.1), 04:50:15/00:03:03, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
```

```
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 04:50:15/00:03:29

(*, 224.0.1.40), 1d05h/00:02:01, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d05h/00:02:01

(*, 225.0.0.102), 1d05h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:00

(172.16.254.6, 225.0.0.102), 1d05h/00:02:05, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:00

(172.16.254.4, 225.0.0.102), 1d05h/00:02:29, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:03:28

(*, 225.0.0.101), 1d05h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:00

(172.16.254.3, 225.0.0.101), 1d05h/00:01:04, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:00
Leaf-02#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Leaf-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
```

```
    Pkts: 0/0/0    Rate: 0 pps
(*,225.0.0.101) Flags: C HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   2/0/170/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A NS
  Tunnel0, VXLAN Decap Flags: F NS
    Pkts: 0/0/0    Rate: 0 pps
(172.16.254.3,225.0.0.101) Flags: HW
  SW Forwarding: 1/0/150/0, Other: 0/0/0
  HW Forwarding:   12630/0/177/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A
  Tunnel0, VXLAN Decap Flags: F NS
    Pkts: 0/0/1    Rate: 0 pps
(*,225.0.0.102) Flags: C HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   1/0/224/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A NS
  Tunnel0, VXLAN Decap Flags: F NS
    Pkts: 0/0/0    Rate: 0 pps
(172.16.254.4,225.0.0.102) Flags: HW
  SW Forwarding: 2/0/163/0, Other: 3/1/2
  HW Forwarding:   9373/0/164/0, Other: 0/0/0
  Null0 Flags: A
  GigabitEthernet1/0/2 Flags: F NS
    Pkts: 0/0/1    Rate: 0 pps
(172.16.254.6,225.0.0.102) Flags: HW
  SW Forwarding: 1/0/206/0, Other: 0/0/0
  HW Forwarding:   3825/0/163/0, Other: 0/0/0
 GigabitEthernet1/0/2 Flags: A
  Tunnel0, VXLAN Decap Flags: F NS
    Pkts: 0/0/1    Rate: 0 pps
(*,232.0.0.0/8) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,239.1.1.1) Flags: C HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   9/0/168/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A NS
  Tunnel0, VXLAN Decap Flags: F NS
    Pkts: 0/0/0    Rate: 0 pps
(172.16.254.4,239.1.1.1) Flags: HW
  SW Forwarding: 1/0/150/0, Other: 24/20/4
  HW Forwarding:   8667/0/156/0, Other: 0/0/0
  Null0 Flags: A
  GigabitEthernet1/0/2 Flags: F NS
    Pkts: 0/0/1    Rate: 0 pps
(172.16.254.6,239.1.1.1) Flags: HW
  SW Forwarding: 2/0/150/0, Other: 0/0/0
  HW Forwarding:   1781/0/168/0, Other: 0/0/0
  GigabitEthernet1/0/2 Flags: A
  Tunnel0, VXLAN Decap Flags: F NS
    Pkts: 0/0/2    Rate: 0 pps
Leaf-02#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 464

Outputs to Verify the Configuration on Border VTEP

The following example shows the output for the **show nve peers** command on Border VTEP:

```
Border# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 1d06h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 1d06h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 1d06h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 1d06h
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  1d06h
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  1d05h
Border#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on Border VTEP:

```
Border# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            6          10101    1d06h
nve1      10102    172.16.254.4            7          10102    1d05h
Border#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Border VTEP:

```
Border# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.6, local AS number 65001
BGP table version is 85, main routing table version 85
5 network entries using 1960 bytes of memory
7 path entries using 1120 bytes of memory
5/5 BGP path/bestpath attribute entries using 1560 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
20 BGP extended community entries using 2706 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7530 total bytes of memory
BGP activity 137/41 prefixes, 272/148 paths, scan interval 60 secs
5 networks peaked at 15:42:39 Aug 6 2020 UTC (1d02h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    2158    2031       85    0    0 1d06h           2
172.16.255.2    4        65001    2157    2025       85    0    0 1d06h           2
Border#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **group-map** command on Border VTEP:

```
Border# show ip pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Vl901,::FFFF:172.16.254.3
    Info source: Static
    Uptime: 1d06h, Groups: 1
Border#
```

The following example shows the output for the **show ip routing vrf** *vrf-name* command on Border VTEP:

```
Border# show ip routing vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "bgp 65001", distance 200, metric 0, type internal
  Redistributing via ospf 1
  Route count is 1/1, share count 0
  Routing paths:
    172.16.254.3%default, Vlan901%default
      From AC10:FF01::
      opaque_ptr 0x7FEF699AEC28
      Last updated 04:34:38 ago
Border#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on Border VTEP:

```
Border# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan102
   1d05h     00:02:29
Border#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on Border VTEP:

```
Border# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 1d05h/00:02:52, RP FC00:2:255::255, flags: SCg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.3
  Immediate Outgoing interface list:
    Vlan102, Null, 1d05h/never
    Vlan2001, Forward, 04:34:39/00:02:52

(FC00:1:102::12, FF06:1::1), 00:15:05/00:02:32, flags: STgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.4
  Immediate Outgoing interface list:
    Vlan2001, Forward, 00:15:03/00:02:32
  Inherited Outgoing interface list:
    Vlan102, Null, 1d05h/never

(FC00:2:255::1, FF06:1::1), 00:14:13/00:02:52, RP FC00:2:255::255, flags: SPR
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.3
  Immediate Outgoing interface list:
    Vlan2001, Null, 00:14:13/00:02:52
  Inherited Outgoing interface list:
    Vlan102, Null, 1d05h/never
```

```
(FC00:2:255::1, FF06:1::1), 00:14:12/never, flags: STGq
  Incoming interface: Vlan2001
  RPF nbr: FE80::A2B4:39FF:FE21:9183
  Immediate Outgoing interface list:
    Vlan901, Forward, 00:14:12/never
  Inherited Outgoing interface list:
    Vlan102, Null, 1d05h/never
Border#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on Border VTEP:

```
Border# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF00::/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF00::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF02::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 12/12/0
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 4/4/0
   HW Forwarding:   7/0/122/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan2001 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:102::12,FF06:1::1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   450/0/125/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan2001 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:2:255::1,FF06:1::1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 2/1/1
   HW Forwarding:   423/0/122/0, Other: 0/0/0
   Vlan2001 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (*,FF10::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF12::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF20::/15) Flags: HW
```

```
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF22::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF30::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF32::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF33::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF34::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF35::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF36::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF37::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF38::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF39::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3A::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3B::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3C::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3D::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3E::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3F::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF40::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF42::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF50::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF52::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF60::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF62::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF70::/15) Flags: HW
```

```
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF72::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF80::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF82::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FF90::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FF92::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFA0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFA2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFB0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFB2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFC0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFC2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFD0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFD2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFE0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFE2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
  (*,FFF0::/15) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,FFF2::/16) Flags:
    SW Forwarding: 0/0/0/0, Other: 0/0/0
Border#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Border VTEP:

```
Border# show bgp ipv6 mvpn all
BGP table version is 85, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                      172.16.255.4             0    100      0 ?
```

```
 *>i                    172.16.255.4             0    100      0 ?
 *>   [5][1:1][FC00:2:255::1][FF06:1::1]/42
                        ::                                    32768 ?
 *>   [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                        ::                                    32768 ?
 * i  [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                        172.16.255.3             0    100      0 ?
 *>i                    172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                        ::                                    32768 ?
Border#
```

The following example shows the output for the **show ip mroute** command on Border VTEP:

```
Border# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d06h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d06h/00:00:08

(172.16.254.6, 239.1.1.1), 00:41:35/00:02:45, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:41:35/00:03:11

(172.16.254.4, 239.1.1.1), 04:57:56/00:02:37, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 04:57:56/00:02:03

(*, 224.0.1.40), 1d06h/00:02:10, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d06h/00:02:10

(*, 225.0.0.102), 1d06h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d06h/00:00:08

(172.16.254.4, 225.0.0.102), 1d06h/00:01:56, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
```

```
    Tunnel0, Forward/Sparse, 1d06h/00:00:08

(172.16.254.6, 225.0.0.102), 1d06h/00:02:16, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
   GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:21, A

(*, 225.0.0.101), 1d06h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d06h/00:00:08

(172.16.254.3, 225.0.0.101), 1d06h/00:02:00, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d06h/00:00:20
Border#
```

The following example shows the output for the **show ip mfib** command on Border VTEP:

```
Border# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   8/0/146/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 7/0/125/0, Other: 0/0/0
   HW Forwarding:   12768/0/177/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/7    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/172/0, Other: 0/0/0
```

```
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,225.0.0.102) Flags: HW
    SW Forwarding: 1/0/154/0, Other: 0/0/0
    HW Forwarding:   9363/0/176/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 17/0/174/0, Other: 10/9/1
    HW Forwarding:   3858/0/151/0, Other: 0/0/0
    Null0 Flags: A
    GigabitEthernet1/0/2 Flags: F
      Pkts: 0/0/16    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   10/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 1/0/150/0, Other: 0/0/0
    HW Forwarding:   8909/0/167/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 2/0/150/0, Other: 2/2/0
    HW Forwarding:   2018/0/156/0, Other: 0/0/0
    Null0 Flags: A
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
Border#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 464

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Spine Switch 1:

```
Spine-01# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 78, main routing table version 78
5 network entries using 1960 bytes of memory
13 path entries using 2080 bytes of memory
4/4 BGP path/bestpath attribute entries using 1216 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2356 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7756 total bytes of memory
BGP activity 270/220 prefixes, 3041/2934 paths, scan interval 60 secs
5 networks peaked at 15:48:28 Aug 6 2020 UTC (1d02h ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
```

```
172.16.255.3    4        65001    2020    2157       78    0    0 1d06h           2
172.16.255.4    4        65001    2030    2154       78    0    0 1d06h           3
172.16.255.6    4        65001    2033    2160       78    0    0 1d06h           3
Spine-01#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 1:

```
Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Spine Switch 1:

```
Spine-01# show bgp ipv6 mvpn all
BGP table version is 78, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
                     172.16.255.6             0    100      0 ?
 *>i                  172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                     172.16.255.6             0    100      0 ?
 * i                  172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                     172.16.255.4             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 * i  [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                     172.16.255.6             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Spine-01#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 1:

```
Spine-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
```

```
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 04:59:49/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 239.1.1.1), 00:43:26/00:02:24, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(172.16.254.4, 239.1.1.1), 04:57:47/00:01:01, flags: PTA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.14.4
  Outgoing interface list: Null

(*, 224.0.1.40), 1w0d/00:02:32, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1w0d/00:02:32

(*, 225.0.0.102), 1w0d/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.102), 02:11:00/00:01:54, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(172.16.254.4, 225.0.0.102), 1d06h/00:02:27, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.14.4
  Outgoing interface list: Null
Spine-01#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 1:

```
Spine-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 83/83/0
```

```
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,224.0.1.40) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    Loopback2 Flags: F IC NS
      Pkts: 0/0/0    Rate: 0 pps
  (*,225.0.0.102) Flags: C HW
    SW Forwarding: 1/0/206/0, Other: 282/0/282
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
  (172.16.254.4,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/2 Flags: NS
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/4 Flags: NS
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 6/5/1
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   8/0/157/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/4 Flags: NS
Spine-01#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 464

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Spine Switch 2:

```
Spine-02# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 77, main routing table version 77
5 network entries using 1960 bytes of memory
13 path entries using 2080 bytes of memory
4/4 BGP path/bestpath attribute entries using 1216 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2356 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7756 total bytes of memory
BGP activity 301/251 prefixes, 3143/3036 paths, scan interval 60 secs
5 networks peaked at 15:49:16 Aug 6 2020 UTC (1d02h ago)
```

```
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    2021    2155       77    0    0 1d06h           2
172.16.255.4    4        65001    2031    2152       77    0    0 1d06h           3
172.16.255.6    4        65001    2029    2161       77    0    0 1d06h           3
Spine-02#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 2:

```
Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Spine Switch 2:

```
Spine-02# show bgp ipv6 mvpn all
BGP table version is 77, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
                     172.16.255.6             0    100      0 ?
 *>i                  172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                     172.16.255.6             0    100      0 ?
 * i                  172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                     172.16.255.4             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 * i  [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                     172.16.255.6             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Spine-02#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 2:

```
Spine-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
```

```
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 3d13h/00:03:01, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:01
    GigabitEthernet1/0/1, Forward/Sparse, 1d06h/00:02:41
    GigabitEthernet1/0/4, Forward/Sparse, 1d06h/00:02:43

(172.16.254.6, 239.1.1.1), 00:44:52/00:02:29, flags: T
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:44:52/00:02:58
    GigabitEthernet1/0/2, Forward/Sparse, 00:44:52/00:03:02

(172.16.254.4, 239.1.1.1), 05:01:13/00:02:28, flags: T
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 05:01:13/00:03:22
    GigabitEthernet1/0/1, Forward/Sparse, 05:01:13/00:02:56

(*, 224.0.1.40), 1w0d/00:03:12, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:12
    GigabitEthernet1/0/1, Forward/Sparse, 1d06h/00:02:54
    GigabitEthernet1/0/4, Forward/Sparse, 1d06h/00:02:44
    Loopback2, Forward/Sparse, 1w0d/00:02:32

(*, 225.0.0.102), 1w0d/00:03:26, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:21
    GigabitEthernet1/0/1, Forward/Sparse, 1d06h/00:03:26
    GigabitEthernet1/0/4, Forward/Sparse, 1d06h/00:02:56

(172.16.254.4, 225.0.0.102), 1d06h/00:02:18, flags: MT
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d06h/00:03:15
    GigabitEthernet1/0/1, Forward/Sparse, 1d06h/00:03:26

(172.16.254.6, 225.0.0.102), 1d06h/00:02:40, flags: MT
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:28
    GigabitEthernet1/0/1, Forward/Sparse, 1d06h/00:03:28

(*, 225.0.0.101), 3d13h/00:03:13, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:02:59
    GigabitEthernet1/0/1, Forward/Sparse, 1d06h/00:02:53
    GigabitEthernet1/0/4, Forward/Sparse, 1d06h/00:03:13
```

```
(172.16.254.3, 225.0.0.101), 1d06h/00:03:09, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:27
    GigabitEthernet1/0/4, Forward/Sparse, 1d06h/00:03:13
Spine-02#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 2:

```
Spine-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 9/0/112/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 4/0/132/0, Other: 0/0/0
   HW Forwarding:   12790/0/177/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 27/0/101/0, Other: 0/0/0
```

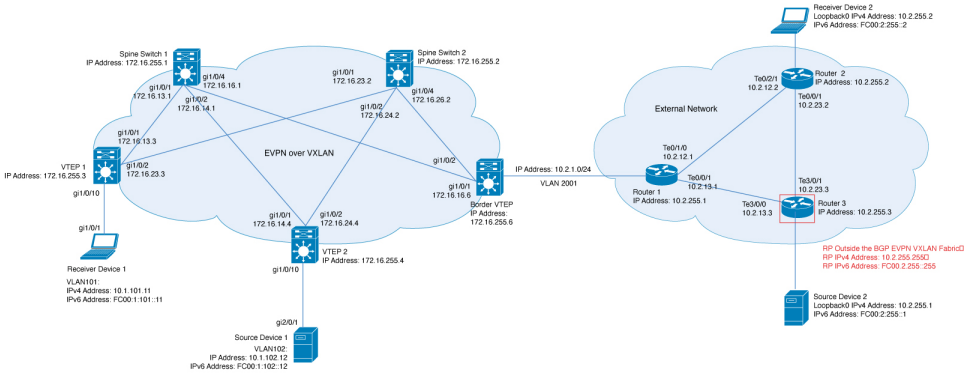
```
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   9381/0/176/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   3853/0/163/0, Other: 0/0/0
    GigabitEthernet1/0/4 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 10/0/150/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/4    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/4    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/4    Rate: 0 pps
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   9007/0/167/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   2111/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/4 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
Spine-02#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric, on page 464

Example: Configuring TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

This example shows how to configure and verify Layer 3 TRM with PIM-SM for IPv4 multicast traffic when the RP is outside the BGP EVPN VXLAN fabric. The example uses the following topology:

Figure 44: TRM with PIM-SM when the RP is Outside the BGP EVPN VXLAN Fabric



The topology shows an EVPN VXLAN network, with two spine switches and three VTEPs, connected to an external network with three routers. Router 3 in the external network acts as the RP in this topology and Border VTEP connects the fabric to the external network through Router 1. The IPv4 multicast group is 226.1.1.1 for this topology. The following tables provide sample configurations for the devices in this topology:

Table 47: Configuring VTEP 1, Border VTEP, and VTEP 2 to Configure TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

VTEP 1	Border VTEP	VTEP 2
Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 !	Border# **show running-config** hostname Border ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 ! vlan 2001 ! interface Loopback0 ip address 172.16.255.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.6 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 !	Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-famil ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 !

VTEP 1	Border VTEP	VTEP 2
interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute connected redistribute static exit-address-family !	interface GigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/5 switchport trunk allowed vlan 2001 switchport mode trunk ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode no autostate ! interface Vlan2001 vrf forwarding green ip address 10.2.1.1 255.255.255.0 ip mtu 1500 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 2 vrf green redistribute bgp 65001	interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 102 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ip pim sparse-mode ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ip pim sparse-mode ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute connected redistribute static exit-address-family !

VTEP 1	Border VTEP	VTEP 2
address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! end ! Leaf-01#	! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static redistribute ospf 2 match internal external 1 external 2 exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! end ! Border#	address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! end ! Leaf-02#

Table 48: Configuring Spine Switch 1 and Spine Switch 2 to Configure TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

Spine Switch 1

```
Spine-01# show running-config
hostname Spine-01
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.14.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/4
no switchport
ip address 172.16.16.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.1
!
router bgp 65001
bgp router-id 172.16.255.1
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0
!
```

Spine Switch 2

```
Spine-02# show running-config
hostname Spine-02
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.23.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.24.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/4
no switchport
ip address 172.16.26.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.2
!
router bgp 65001
bgp router-id 172.16.255.2
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0
!
```

Spine Switch 1	Spine Switch 2
address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-01#	address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-02#

Table 49: Configuring Router 1, Router 2, and Router 3 to Configure TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

Router 1	Router 2	Router 3
R1# **show running-config** hostname R1 ! ip multicast-routing distributed ! interface Loopback0 ip address 10.2.255.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet0/0/1.2001 encapsulation dot1Q 2001 ip address 10.2.1.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 router ospf 1 router-id 10.2.255.1 ! ip pim rp-address 10.2.255.255 ! end ! R1#	R2# **show running-config** hostname R2 ! ip multicast-routing distributed ! interface Loopback0 ip address 10.2.255.2 255.255.255.255 ip pim sparse-mode ip igmp join-group 226.1.1.1 ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 10.2.255.2 ! ip pim rp-address 10.2.255.255 ! end ! R2#	R3# **show running-config** hostname R3 ! ip multicast-routing distributed ! interface Loopback0 ip address 10.2.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback255 ip address 10.2.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 10.2.255.3 ! ip pim rp-address 10.2.255.255 ! end ! R3#

Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

The following sections provide sample outputs for **show** commands to verify TRM with PIM-SM on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 502
- Outputs to Verify the Configuration on VTEP 2, on page 507
- Outputs to Verify the Configuration on Border VTEP, on page 512
- Outputs to Verify the Configuration on Spine Switch 1, on page 518
- Outputs to Verify the Configuration on Spine Switch 2, on page 521
- Outputs to Verify the Configuration on Router 3 (RP Outside the BGP EVPN VXLAN Fabric), on page 525

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Leaf-01# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 16:44:02
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 16:41:00
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 16:44:02
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 16:41:00
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  16:23:05
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  16:44:02
Leaf-01#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
Leaf-01# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10102    172.16.254.4            7          10102    16:23:06
nve1      10102    172.16.254.6            5          10102    16:44:02
Leaf-01#
```

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on VTEP 1:

```
Leaf-01# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 58, main routing table version 58
6 network entries using 1824 bytes of memory
8 path entries using 1088 bytes of memory
3/3 BGP path/bestpath attribute entries using 936 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2372 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6404 total bytes of memory
BGP activity 117/25 prefixes, 240/113 paths, scan interval 60 secs
6 networks peaked at 12:17:52 Aug 6 2020 UTC (16:27:28.286 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1217    1126       58    0    0 16:44:58        2
172.16.255.2    4        65001    1213    1121       58    0    0 16:44:53        2
Leaf-01#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **rp mapping** command on VTEP 1:

```
Leaf-01# show ip pim vrf green rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
Leaf-01#
```

The following example shows the output for the **show ip routing vrf** command on VTEP 1:

```
Leaf-01# show ip routing vrf green 10.2.255.255
Routing Table: green
Routing entry for 10.2.255.255/32
  Known via "bgp 65001", distance 200, metric 3, type internal
  Last update from 172.16.254.6 on Vlan901, 16:17:01 ago
  Routing Descriptor Blocks:
  * 172.16.254.6 (default), from 172.16.255.1, 16:17:01 ago, via Vlan901
      opaque_ptr 0x7FBB8620D990
      Route metric is 3, traffic share count is 1
      AS Hops 0
      MPLS label: none
Leaf-01#
```

The following example shows the output for the **show ip igmp vrf** *vrf-name* **groups** command on VTEP 1:

```
Leaf-01# show ip igmp vrf green groups
IGMP Connected Group Membership
Group Address    Interface                Uptime    Expires   Last Reporter   Group Accounted
226.1.1.1        Vlan101                  13:03:08  00:02:13  10.1.101.11
224.0.1.40       Loopback901              16:45:17  00:02:50  10.1.255.1
Leaf-01#
```

The following example shows the output for the **show ip mroute vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 13:03:08/stopped, RP 10.2.255.255, flags: SJCg
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list:
    Vlan101, Forward/Sparse, 13:03:08/00:02:13

(10.2.255.1, 226.1.1.1), 00:08:23/00:02:54, flags: TgQ
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list:
    Vlan101, Forward/Sparse, 00:08:23/00:02:13

(10.1.102.12, 226.1.1.1), 00:08:48/00:02:24, flags: TgQ
  Incoming interface: Vlan901, RPF nbr 172.16.254.4
  Outgoing interface list:
    Vlan101, Forward/Sparse, 00:08:48/00:02:13

(*, 224.0.1.40), 16:45:17/00:02:50, RP 10.2.255.255, flags: SJCLg
```

```
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
Leaf-01#
```

The following example shows the output for the **show ip mfib vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ip mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 2/2/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Loopback901 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,226.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan101 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (10.1.102.12,226.1.1.1) Flags: HW
   SW Forwarding: 5/0/100/0, Other: 0/0/0
   HW Forwarding:   523/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
Leaf-01#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on VTEP 1:

```
Leaf-01# show bgp ipv4 mvpn all
BGP table version is 60, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 * i                    172.16.255.4             0    100      0 ?
 *>   [5][1:1][10.2.255.1][226.1.1.1]/18
                      0.0.0.0                            32768 ?
```

```
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                     172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                     172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
 *>i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                     172.16.255.3             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                     0.0.0.0                            32768 ?
Leaf-01#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Leaf-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 16:45:08/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 16:45:08/00:00:45

(172.16.254.4, 239.1.1.1), 00:08:47/00:01:59, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:08:47/00:00:12

(172.16.254.6, 239.1.1.1), 00:08:49/00:02:00, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:08:49/00:00:10

(*, 224.0.1.40), 16:45:17/00:02:46, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
Leaf-01#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Leaf-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
```

```
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/114/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 13/0/127/0, Other: 2/2/0
   HW Forwarding:   7870/0/164/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   5222/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   2137/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   11/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 4/0/150/0, Other: 0/0/0
```

```
   HW Forwarding:   518/0/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   498/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
Leaf-01#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 501

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Leaf-02# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 16:56:53
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 16:56:53
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 16:56:53
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 16:56:53
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  16:56:53
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  16:56:53
Leaf-02#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
Leaf-02# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            6          10101    16:56:54
nve1      10102    172.16.254.6            5          10102    16:56:54
Leaf-02#
```

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on VTEP 2:

```
Leaf-02# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 62, main routing table version 62
7 network entries using 2128 bytes of memory
9 path entries using 1224 bytes of memory
4/4 BGP path/bestpath attribute entries using 1248 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2372 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7156 total bytes of memory
BGP activity 121/28 prefixes, 202/77 paths, scan interval 60 secs
9 networks peaked at 12:22:24 Aug 6 2020 UTC (16:43:21.423 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1229    1151       62    0    0 16:57:50        2
```

```
172.16.255.2    4        65001    1227    1152       62    0    0 16:57:51        2
Leaf-02#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **rp mapping** command on VTEP 2:

```
Leaf-02# show ip pim vrf green rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
Leaf-02#
```

The following example shows the output for the **show ip routing vrf** command on VTEP 2:

```
Leaf-02# show ip routing vrf green 10.2.255.255
Routing Table: green
Routing entry for 10.2.255.255/32
  Known via "bgp 65001", distance 200, metric 3, type internal
  Last update from 172.16.254.6 on Vlan901, 16:56:55 ago
  Routing Descriptor Blocks:
  * 172.16.254.6 (default), from 172.16.255.1, 16:56:55 ago, via Vlan901
      opaque_ptr 0x7F65B8B9E4B0
      Route metric is 3, traffic share count is 1
      AS Hops 0
      MPLS label: none
Leaf-02#
```

The following example shows the output for the **show ip igmp vrf** *vrf-name* **groups** command on VTEP 2:

```
Leaf-02# show ip igmp vrf green groups
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter   Group Accounted
226.1.1.1        Vlan102                  16:58:00  00:02:11  10.1.102.12
224.0.1.40       Vlan901                  16:58:37  00:02:33  172.16.254.4
Leaf-02#
```

The following example shows the output for the **show ip mroute vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
```

```
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 16:58:00/stopped, RP 10.2.255.255, flags: SJCFg
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list:
    Vlan102, Forward/Sparse, 16:58:00/00:02:11

(10.2.255.1, 226.1.1.1), 00:24:16/00:02:40, flags: JTgQ
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list:
    Vlan102, Forward/Sparse, 00:24:16/00:02:11

(10.1.102.12, 226.1.1.1), 00:24:41/00:02:09, flags: FTGqx
  Incoming interface: Vlan102, RPF nbr 0.0.0.0, Registering
  Outgoing interface list:
    Vlan901, Forward/Sparse, 00:24:41/stopped

(*, 224.0.1.40), 16:58:37/00:02:33, RP 10.2.255.255, flags: SJPCLgx
  Incoming interface: Vlan901, RPF nbr 172.16.254.6
  Outgoing interface list: Null
Leaf-02#
```

The following example shows the output for the **show ip mfib vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ip mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A IC NS
 (*,226.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   3/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan102 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (10.1.102.12,226.1.1.1) Flags: HW
   SW Forwarding: 739/0/100/0, Other: 2/2/0
   HW Forwarding:   736/0/118/0, Other: 0/0/0
   Vlan102 Flags: A
   Tunnel5 Flags: F
     Pkts: 0/0/739    Rate: 0 pps
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
```

```
     Pkts: 0/0/739    Rate: 0 pps
Leaf-02#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on VTEP 2:

```
Leaf-02# show bgp ipv4 mvpn all
BGP table version is 62, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>   [5][1:1][10.1.102.12][226.1.1.1]/18
                       0.0.0.0                                32768 ?
 *>i  [5][1:1][10.2.255.1][226.1.1.1]/18
                       172.16.255.6             0    100      0 ?
 * i                    172.16.255.6             0    100      0 ?
 *>   [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                       0.0.0.0                                32768 ?
 *>   [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                       0.0.0.0                                32768 ?
 *>i  [7][1:1][65001][10.1.102.12/32][226.1.1.1/32]/22
                       172.16.255.3             0    100      0 ?
 *>   [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                       0.0.0.0                                32768 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                       172.16.255.3             0    100      0 ?
 * i                    172.16.255.3             0    100      0 ?
Leaf-02#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Leaf-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 16:58:28/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 16:58:28/00:02:25
```

```
(172.16.254.6, 239.1.1.1), 00:24:42/00:00:58, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:24:42/00:02:17

(172.16.254.4, 239.1.1.1), 00:24:42/00:03:28, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:24:42/00:03:23, A

(*, 224.0.1.40), 16:58:37/00:02:26, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 16:58:36/00:02:26
Leaf-02#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Leaf-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/170/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   7870/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/224/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
```

```
    Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 2/0/163/0, Other: 3/1/2
   HW Forwarding:   5353/0/164/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/206/0, Other: 0/0/0
   HW Forwarding:   2165/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   5/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 1495/1491/4
   HW Forwarding:   742/0/156/0, Other: 0/0/0
   Null0 Flags: A NS
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   1460/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
Leaf-02#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 501

Outputs to Verify the Configuration on Border VTEP

The following example shows the output for the **show nve peers** command on Border VTEP:

```
Border# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 17:09:20
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 17:06:19
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 17:09:20
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 17:06:19
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  17:09:20
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  16:48:24
Border#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on Border VTEP:

Border VTEP

```
Border# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
```

```
--------- -------- ------------------------  ---------- -------- --------
nve1      10101    172.16.254.3              6          10101    17:09:21
nve1      10102    172.16.254.4              7          10102    16:48:24
Border#
```

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on Border VTEP:

```
Border# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.6, local AS number 65001
BGP table version is 60, main routing table version 60
6 network entries using 1824 bytes of memory
10 path entries using 1360 bytes of memory
4/4 BGP path/bestpath attribute entries using 1248 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
19 BGP extended community entries using 2682 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7298 total bytes of memory
BGP activity 116/24 prefixes, 232/112 paths, scan interval 60 secs
8 networks peaked at 12:14:22 Aug 6 2020 UTC (16:52:46.174 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1246    1165       60    0    0 17:13:17        4
172.16.255.2    4        65001    1247    1161       60    0    0 17:13:14        4
Border#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **rp mapping** command on Border VTEP:

```
Border# show ip pim vrf green rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
Border#
```

The following example shows the output for the **show ip routing vrf** *vrf-name* command on Border VTEP:

```
Border# show ip routing vrf green 10.2.255.255
Routing Table: green
Routing entry for 10.2.255.255/32
  Known via "ospf 2", distance 110, metric 3, type intra area
  Redistributing via bgp 65001
  Advertised by bgp 65001 match internal external 1 & 2
  Last update from 10.2.1.2 on Vlan2001, 17:12:42 ago
  Routing Descriptor Blocks:
  * 10.2.1.2, from 10.2.255.3, 17:12:42 ago, via Vlan2001
      Route metric is 3, traffic share count is 1
Border#
```

The following example shows the output for the **show ip igmp vrf** *vrf-name* **groups** command on Border VTEP:

```
Border# show ip igmp vrf green groups
IGMP Connected Group Membership
Group Address    Interface              Uptime    Expires   Last Reporter   Group Accounted
```

```
224.0.1.40       Vlan901                 17:14:13  00:02:51  172.16.254.6
Border#
```

The following example shows the output for the **show ip mroute vrf** *vrf-name* command on Border VTEP:

```
Border# show ip mroute vrf green
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 17:06:19/stopped, RP 10.2.255.255, flags: SJGx
  Incoming interface: Vlan2001, RPF nbr 10.2.1.2
  Outgoing interface list:
    Vlan901, Forward/Sparse, 17:06:19/stopped

(10.2.255.1, 226.1.1.1), 00:33:41/00:01:22, flags: TGqx
  Incoming interface: Vlan2001, RPF nbr 10.2.1.2
  Outgoing interface list:
    Vlan901, Forward/Sparse, 00:33:41/stopped

(10.1.102.12, 226.1.1.1), 00:34:06/00:03:14, flags: Tgx
  Incoming interface: Vlan901, RPF nbr 172.16.254.4
  Outgoing interface list:
    Vlan2001, Forward/Sparse, 00:34:06/00:02:52, A

(*, 224.0.1.40), 17:14:13/00:02:51, RP 10.2.255.255, flags: SJCLGx
  Incoming interface: Vlan2001, RPF nbr 10.2.1.2
  Outgoing interface list:
    Vlan901, Forward/Sparse, 17:14:12/00:02:51
Border#
```

The following example shows the output for the **show ip mfib vrf** *vrf-name* command on Border VTEP:

```
Border# show ip mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
```

```
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Vlan2001 Flags: A NS
   Vlan901, VXLAN Decap Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,226.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   6/0/122/0, Other: 0/0/0
   Vlan2001 Flags: A NS
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (10.1.102.12,226.1.1.1) Flags: HW
   SW Forwarding: 4/0/100/0, Other: 1/1/0
   HW Forwarding:   2096/1/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan2001 Flags: F
     Pkts: 0/0/4    Rate: 0 pps
 (10.2.255.1,226.1.1.1) Flags: HW
   SW Forwarding: 1/0/100/0, Other: 0/0/0
   HW Forwarding:   2072/1/122/0, Other: 0/0/0
   Vlan2001 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/1    Rate: 0 pps
Border#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on Border VTEP:

```
Border# show bgp ipv4 mvpn all
BGP table version is 60, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 * i                    172.16.255.4             0    100      0 ?
 *>   [5][1:1][10.2.255.1][226.1.1.1]/18
                      0.0.0.0                            32768 ?
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                      172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                      172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
 *>i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                      172.16.255.3             0    100      0 ?
```

```
 * i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                      0.0.0.0                                32768 ?
Border#
```

The following example shows the output for the **show ip mroute** command on Border VTEP:

```
Border# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 17:14:04/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 17:14:04/00:01:48

(172.16.254.4, 239.1.1.1), 00:34:05/00:02:44, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:34:05/00:01:54

(172.16.254.6, 239.1.1.1), 00:34:07/00:03:12, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:34:07/00:02:52, A

(*, 224.0.1.40), 17:14:13/00:02:47, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 17:14:12/00:02:47
Border#
```

The following example shows the output for the **show ip mfib** command on Border VTEP:

```
Border# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
```

```
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   8/0/146/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 7/0/125/0, Other: 0/0/0
   HW Forwarding:   8010/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/7    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   5353/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 11/0/167/0, Other: 7/6/1
   HW Forwarding:   2207/0/151/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/10    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   8/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 4/0/150/0, Other: 0/0/0
   HW Forwarding:   2032/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 2/0/150/0, Other: 4/4/0
```

```
   HW Forwarding:   2015/1/156/1, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/1    Rate: 0 pps
Border#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 501

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on Spine Switch 1:

```
Spine-01# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 169, main routing table version 169
6 network entries using 1824 bytes of memory
16 path entries using 2176 bytes of memory
3/3 BGP path/bestpath attribute entries using 912 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
16 BGP extended community entries using 2332 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7388 total bytes of memory
BGP activity 250/203 prefixes, 2984/2883 paths, scan interval 60 secs
8 networks peaked at 12:20:11 Aug 6 2020 UTC (16:59:40.011 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    1161    1252      169    0    0 17:17:09        4
172.16.255.4    4        65001    1169    1247      169    0    0 17:14:09        4
172.16.255.6    4        65001    1172    1253      169    0    0 17:20:10        2
Spine-01#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 1:

```
Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-01#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on Spine Switch 1:

```
Spine-01# show bgp ipv4 mvpn all
BGP table version is 169, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
```

```
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [5][1:1][10.2.255.1][226.1.1.1]/18
                     172.16.255.6             0    100      0 ?
 *>i                  172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 *>i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                     172.16.255.3             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 * i                  172.16.255.4             0    100      0 ?
 *>i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                     172.16.255.3             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 * i                  172.16.255.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                     172.16.255.3             0    100      0 ?
 * i                  172.16.255.3             0    100      0 ?
 * i                  172.16.255.6             0    100      0 ?
Spine-01#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 1:

```
Spine-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 00:42:45/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 239.1.1.1), 00:42:22/00:02:37, flags: PTA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(172.16.254.4, 239.1.1.1), 00:42:45/00:02:28, flags: PTA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.14.4
  Outgoing interface list: Null

(*, 224.0.1.40), 1w0d/00:02:18, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1w0d/00:02:18
```

```
(*, 225.0.0.102), 6d19h/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.102), 05:29:52/00:02:22, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(172.16.254.4, 225.0.0.102), 17:12:35/00:02:03, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.14.4
  Outgoing interface list: Null
Spine-01#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 1:

```
Spine-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 82/82/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 1/0/206/0, Other: 279/0/279
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: NS
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/4 Flags: NS
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 1/0/1
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
```

```
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1224/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/4 Flags: A NS
Spine-01#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 501

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp ipv4 mvpn all summary** command on Spine Switch 2:

```
Spine-02# show bgp ipv4 mvpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 131, main routing table version 131
6 network entries using 1824 bytes of memory
16 path entries using 2176 bytes of memory
3/3 BGP path/bestpath attribute entries using 912 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
16 BGP extended community entries using 2332 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7388 total bytes of memory
BGP activity 283/236 prefixes, 3089/2988 paths, scan interval 60 secs
8 networks peaked at 12:20:59 Aug 6 2020 UTC (17:02:43.558 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    1160    1252      131    0    0 17:20:09        4
172.16.255.4    4        65001    1173    1249      131    0    0 17:17:14        4
172.16.255.6    4        65001    1172    1258      131    0    0 17:23:12        2
Spine-02#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 2:

```
Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-02#
```

The following example shows the output for the **show bgp ipv4 mvpn all** command on Spine Switch 2:

```
Spine-02# show bgp ipv4 mvpn all
BGP table version is 131, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][10.1.102.12][226.1.1.1]/18
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.4             0    100      0 ?
 * i  [5][1:1][10.2.255.1][226.1.1.1]/18
                      172.16.255.6             0    100      0 ?
 *>i                   172.16.255.6             0    100      0 ?
 * i  [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
                      172.16.255.4             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 *>i  [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
                      172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
 *>i  [7][1:1][65001][10.2.255.1/32][226.1.1.1/32]/22
                      172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [7][172.16.254.4:102][65001][10.1.102.12/32][226.1.1.1/32]/22
                      172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
Spine-02#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 2:

```
Spine-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 3d00h/00:03:23, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 17:17:14/00:03:23
    GigabitEthernet1/0/1, Forward/Sparse, 17:20:16/00:03:17
    GigabitEthernet1/0/4, Forward/Sparse, 17:23:12/00:02:52

(172.16.254.4, 239.1.1.1), 00:44:04/00:01:34, flags: T
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:44:04/00:02:52
    GigabitEthernet1/0/4, Forward/Sparse, 00:44:04/00:03:17
```

```
(172.16.254.6, 239.1.1.1), 00:44:04/00:01:32, flags: T
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:44:04/00:03:17
    GigabitEthernet1/0/2, Forward/Sparse, 00:44:04/00:03:23

(*, 224.0.1.40), 1w0d/00:03:22, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 17:17:14/00:02:46
    GigabitEthernet1/0/1, Forward/Sparse, 17:20:16/00:03:22
    GigabitEthernet1/0/4, Forward/Sparse, 17:23:12/00:03:13
    Loopback2, Forward/Sparse, 1w0d/00:02:33

(*, 225.0.0.102), 1w0d/00:03:29, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 17:17:14/00:03:04
    GigabitEthernet1/0/1, Forward/Sparse, 17:20:16/00:03:29
    GigabitEthernet1/0/4, Forward/Sparse, 17:23:12/00:02:36

(172.16.254.4, 225.0.0.102), 17:17:08/00:02:44, flags: MT
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 17:17:08/00:02:37
    GigabitEthernet1/0/1, Forward/Sparse, 17:17:08/00:03:29

(172.16.254.6, 225.0.0.102), 17:23:14/00:03:21, flags: MT
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 17:17:14/00:03:04
    GigabitEthernet1/0/1, Forward/Sparse, 17:20:16/00:03:29

(*, 225.0.0.101), 3d00h/00:03:10, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 17:17:14/00:03:01
    GigabitEthernet1/0/1, Forward/Sparse, 17:20:16/00:03:10
    GigabitEthernet1/0/4, Forward/Sparse, 17:23:12/00:02:40

(172.16.254.3, 225.0.0.101), 17:19:56/00:02:53, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 17:17:14/00:03:01
    GigabitEthernet1/0/4, Forward/Sparse, 17:19:56/00:03:02
Spine-02#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 2:

```
Spine-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup
```

```
Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 9/0/112/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 4/0/132/0, Other: 0/0/0
   HW Forwarding:   8067/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 27/0/101/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   5404/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2214/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/4 Flags: A NS
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
```

```
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 9/0/150/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
    Tunnel1 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/3    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/3    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/3    Rate: 0 pps
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   2629/1/168/1, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/4 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   2607/1/168/1, Other: 0/0/0
    GigabitEthernet1/0/4 Flags: A
    GigabitEthernet1/0/1 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
Spine-02#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 501

Outputs to Verify the Configuration on Router 3 (RP Outside the BGP EVPN VXLAN Fabric)

The following example shows the output for the **show ip pim rp mapping** command on Router 3:

```
R3# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 10.2.255.255 (?)
R3#
```

The following example shows the output for the **show ip mroute** command on Router 3:

```
R3# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 226.1.1.1), 2d19h/00:03:17, RP 10.2.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet0/0/0, Forward/Sparse, 17:10:20/00:03:17
    GigabitEthernet0/0/1, Forward/Sparse, 2d16h/00:03:11

(10.2.255.1, 226.1.1.1), 00:37:40/00:02:14, flags: PJT
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 10.2.13.1
  Outgoing interface list: Null

(10.1.102.12, 226.1.1.1), 00:38:05/00:02:58, flags: P
  Incoming interface: GigabitEthernet0/0/0, RPF nbr 10.2.13.1
  Outgoing interface list: Null

(*, 224.0.1.40), 2d23h/00:03:27, RP 10.2.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    Loopback0, Forward/Sparse, 2d23h/00:02:11
    GigabitEthernet0/0/1, Forward/Sparse, 2d19h/00:03:26
    GigabitEthernet0/0/0, Forward/Sparse, 2d19h/00:03:27
R3#
```

Return to Verifying TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 501

Example: Configuring TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when RP is Outside the BGP EVPN VXLAN Fabric

This example shows how to configure and verify Layer 3 TRM with PIM-SM for IPv4 and IPv6 multicast traffic when the RP is outside the BGP EVPN VXLAN fabric. The example uses the topology in the TRM with PIM-SM when the RP is Outside the BGP EVPN VXLAN Fabric figure.

The topology shows an EVPN VXLAN network, with two spine switches and three VTEPs, connected to an external network with three routers. Router 3 in the external network acts as the RP in this topology and Border VTEP connects the fabric to the external network through Router 1. The IPv4 multicast group is 226.1.1.1 and the IPv6 multicast group is FF06:1::1 in this topology. The following tables provide sample configurations for the devices in this topology:

Table 50: Configuring VTEP 1, Border VTEP, and VTEP 2 to Configure TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

VTEP 1	Border VTEP	VTEP 2
Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901	Border# **show running-config** hostname Border ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !	Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 mdt auto-discovery vxlan mdt default vxlan 239.1.1.1 mdt overlay use-bgp route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ip multicast-routing vrf green ! ipv6 unicast-routing ipv6 multicast-routing vrf green ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 102 member evpn-instance 102 vni 10102 vlan configuration 901 member vni 50901 !

VTEP 1	Border VTEP	VTEP 2
!	vlan 2001	
interface Loopback0	!	interface Loopback0
ip address 172.16.255.3	interface Loopback0	ip address 172.16.255.4
255.255.255.255	ip address 172.16.255.6	255.255.255.255
ip pim sparse-mode	255.255.255.255	ip pim sparse-mode
ip ospf 1 area 0	ip pim sparse-mode	ip ospf 1 area 0
!	ip ospf 1 area 0	!
interface Loopback1	!	interface Loopback1
ip address 172.16.254.3	interface Loopback1	ip address 172.16.254.4
255.255.255.255	ip address 172.16.254.6	255.255.255.255
ip pim sparse-mode	255.255.255.255	ip pim sparse-mode
ip ospf 1 area 0	ip pim sparse-mode	ip ospf 1 area 0
!	ip ospf 1 area 0	!
interface Loopback901	!	interface Loopback901
vrf forwarding green	interface Loopback901	vrf forwarding green
ip address 10.1.255.1 255.255.255.255	vrf forwarding green	ip address 10.1.255.2 255.255.255.255
ip pim sparse-mode	ip address 10.1.255.4 255.255.255.255	ip pim sparse-mode
ipv6 address FC00:1:255::1/128	ip pim sparse-mode	ipv6 address FC00:1:255::2/128
ipv6 enable	ipv6 address FC00:1:255::4/128	ipv6 enable
!	ipv6 enable	!
interface GigabitEthernet1/0/1	!	interface GigabitEthernet1/0/1
no switchport	interface GigabitEthernet1/0/1	no switchport
ip address 172.16.13.3 255.255.255.0	no switchport	ip address 172.16.14.4 255.255.255.0
ip pim sparse-mode	ip address 172.16.16.6 255.255.255.0	ip pim sparse-mode
ip ospf network point-to-point	ip pim sparse-mode	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf network point-to-point	ip ospf 1 area 0
!	ip ospf 1 area 0	!
interface GigabitEthernet1/0/2	!	interface GigabitEthernet1/0/2
no switchport	interface GigabitEthernet1/0/2	no switchport
ip address 172.16.23.3 255.255.255.0	no switchport	ip address 172.16.24.4 255.255.255.0
ip pim sparse-mode	ip address 172.16.26.6 255.255.255.0	ip pim sparse-mode
ip ospf network point-to-point	ip pim sparse-mode	ip ospf network point-to-point
ip ospf 1 area 0	ip ospf network point-to-point	ip ospf 1 area 0
!	ip ospf 1 area 0	!
interface GigabitEthernet1/0/10	!	interface GigabitEthernet1/0/10
switchport access vlan 101	interface GigabitEthernet1/0/5	switchport access vlan 101
switchport mode access	switchport trunk allowed vlan 2001	switchport mode access
!	switchport mode trunk	!
interface Vlan101	!	interface Vlan101
vrf forwarding green	interface GigabitEthernet1/0/10	vrf forwarding green
ip address 10.1.101.1 255.255.255.0	switchport access vlan 102	ip address 10.1.101.1 255.255.255.0
ip pim sparse-mode	switchport mode access	ip pim sparse-mode
ipv6 address FC00:1:101::1/64	!	ipv6 address FC00:1:101::1/64
ipv6 enable	interface Vlan101	ipv6 enable
!	vrf forwarding green	!
interface Vlan102	ip address 10.1.101.1 255.255.255.0	interface Vlan102
vrf forwarding green	ip pim sparse-mode	vrf forwarding green
ip address 10.1.102.1 255.255.255.0	ipv6 address FC00:1:101::1/64	ip address 10.1.102.1 255.255.255.0
ip pim sparse-mode	ipv6 enable	ip pim sparse-mode
ipv6 address FC00:1:102::1/64	!	ipv6 address FC00:1:102::1/64
ipv6 enable	interface Vlan102	ipv6 enable
!	vrf forwarding green	!
interface Vlan901	ip address 10.1.102.1 255.255.255.0	interface Vlan901
vrf forwarding green	ip pim sparse-mode	vrf forwarding green
ip unnumbered Loopback1	ipv6 address FC00:1:102::1/64	ip unnumbered Loopback1
ip pim sparse-mode	ipv6 enable	ip pim sparse-mode
ipv6 enable	!	ipv6 enable
no autostate		no autostate
!		!

VTEP 1	Border VTEP	VTEP 2
interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute connected redistribute static exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family !	interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable no autostate ! interface Vlan2001 vrf forwarding green ip address 10.2.1.1 255.255.255.0 ip mtu 1500 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ipv6 address FC00:2:1::1/64 ipv6 enable ipv6 mtu 1500 ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospfv3 1 ! address-family ipv6 unicast vrf green redistribute bgp 65001 exit-address-family ! router ospf 2 vrf green redistribute bgp 65001 ! router ospf 1 router-id 172.16.255.6 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0	interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 mcast-group 225.0.0.102 ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute connected redistribute static exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family !

VTEP 1	Border VTEP	VTEP 2
address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end ! Leaf-01#	! address-family ipv4 exit-address-family ! address-family ipv4 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static redistribute ospf 2 match internal external 1 external 2 exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute ospf 1 include-connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ! end ! Border#	address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip pim vrf green rp-address 10.2.255.255 ! ipv6 pim vrf green rp-address FC00:2:255::255 ipv6 pim vrf green register-source Loopback901 ! end ! Leaf-02#

Table 51: Configuring Spine Switch 1 and Spine Switch 2 to Configure TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

Spine Switch 1	Spine Switch 2

Spine Switch 1

```
Spine-01# show running-config
hostname Spine-01
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.14.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/4
no switchport
ip address 172.16.16.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.1
!
router bgp 65001
bgp router-id 172.16.255.1
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0
!
address-family ipv4
exit-address-family
!
```

Spine Switch 2

```
Spine-02# show running-config
hostname Spine-02
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.23.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.24.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/4
no switchport
ip address 172.16.26.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.2
!
router bgp 65001
bgp router-id 172.16.255.2
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.6 remote-as 65001
neighbor 172.16.255.6 update-source Loopback0
!
address-family ipv4
exit-address-family
!
```

Spine Switch 1	Spine Switch 2
address-family ipv4 mvpn neighbor 172.16.255.2 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-01#	address-family ipv4 mvpn neighbor 172.16.255.1 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family ipv6 mvpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.6 activate neighbor 172.16.255.6 send-community both neighbor 172.16.255.6 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip pim ssm default ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end ! Spine-02#

Table 52: Configuring Router 1, Router 2, and Router 3 to Configure TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

Router 1	Router 2	Router 3
R1# **show running-config** hostname R1 ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! interface Loopback0 ip address 10.2.255.1 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ipv6 address FC00:2:255::1/128 ipv6 enable ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:12::1/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:13::1/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface GigabitEthernet0/0/1.2001 encapsulation dot1Q 2001 ip address 10.2.1.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ipv6 address FC00:2:1::2/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0	R2# **show running-config** hostname R2 ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! interface Loopback0 ip address 10.2.255.2 255.255.255.255 ip pim sparse-mode ip igmp join-group 226.1.1.1 ip ospf 1 area 0 ipv6 address FC00:2:255::2/128 ipv6 enable ipv6 mld join-group FF06:1::1 ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.12.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:12::2/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:23::2/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! router ospfv3 1 ! address-family ipv6 unicast exit-address-family ! router ospf 1 router-id 10.2.255.2	R3# **show running-config** hostname R3 ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! interface Loopback0 ip address 10.2.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ipv6 address FC00:2:255::3/128 ipv6 enable ospfv3 1 ipv6 area 0 ! interface Loopback255 ip address 10.2.255.255 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ipv6 address FC00:2:255::255/128 ipv6 enable ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/0 ip address 10.2.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:13::3/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! interface TenGigabitEthernet0/0/1 ip address 10.2.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 cdp enable ipv6 address FC00:2:23::3/64 ipv6 enable ospfv3 network point-to-point ospfv3 1 ipv6 area 0 ! router ospfv3 1 !

Router 1	Router 2	Router 3
! router ospfv3 1 ! address-family ipv6 unicast exit-address-family ! router ospf 1 router-id 10.2.255.1 ! ip pim rp-address 10.2.255.255 ! ipv6 pim rp-address FC00:2:255::255 ! end ! R1#	! ip pim rp-address 10.2.255.255 ! ipv6 pim rp-address FC00:2:255::255 ! end ! R2#	address-family ipv6 unicast exit-address-family ! router ospf 1 router-id 10.2.255.3 ! ip pim rp-address 10.2.255.255 ! ipv6 pim rp-address FC00:2:255::255 ! end ! R3#

Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric

The following sections provide sample outputs for **show** commands to verify TRM with PIM-SM on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 535
- Outputs to Verify the Configuration on VTEP 2, on page 543
- Outputs to Verify the Configuration on Border VTEP, on page 550
- Outputs to Verify the Configuration on Spine Switch 1, on page 558
- Outputs to Verify the Configuration on Spine Switch 2, on page 561
- Outputs to Verify the Configuration on Router 3 (RP Outside the BGP EVPN VXLAN Fabric), on page 565

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peers** command on VTEP 1:

```
Leaf-01# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 1d01h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 1d01h
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 1d01h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 1d01h
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  1d00h
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  1d01h
Leaf-01#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
Leaf-01# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
```

```
--------- -------- ------------------------  ---------- -------- --------
nve1      10102    172.16.254.4              7          10102    1d00h
nve1      10102    172.16.254.6              5          10102    1d01h
Leaf-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on VTEP 1:

```
Leaf-01# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 43, main routing table version 43
5 network entries using 1960 bytes of memory
7 path entries using 1120 bytes of memory
3/3 BGP path/bestpath attribute entries using 936 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2372 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 6572 total bytes of memory
BGP activity 124/30 prefixes, 253/123 paths, scan interval 60 secs
5 networks peaked at 15:46:09 Aug 6 2020 UTC (21:27:07.275 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1796    1688       43    0    0 1d01h           2
172.16.255.2    4        65001    1795    1685       43    0    0 1d01h           2
Leaf-01#
```

The following example shows the output for the **show ipv6 pim vrf** *vrf-name* **group-map** command on VTEP 1:

```
Leaf-01# show ip pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Vl901,::FFFF:172.16.254.6
    Info source: Static
    Uptime: 21:43:02, Groups: 1
Leaf-01#
```

The following example shows the output for the **show ipv6 routing vrf** command on VTEP 1:

```
Leaf-01# show ipv6 routing vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "bgp 65001", distance 200, metric 2, type internal
  Route count is 1/1, share count 0
  Routing paths:
    172.16.254.6%default, Vlan901%default
      From AC10:FF01::
      opaque_ptr 0x7FBB863DE268
      Last updated 1d00h ago
Leaf-01#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 1:

```
Leaf-01# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan101
   21:30:55  00:03:57
Leaf-01#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 21:30:55/00:03:08, RP FC00:2:255::255, flags: SCJg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.6
  Immediate Outgoing interface list:
    Vlan101, Forward, 21:30:55/00:03:08

(FC00:1:102::12, FF06:1::1), 00:01:55/00:01:34, flags: SJTgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.4
  Immediate Outgoing interface list:
    Vlan101, Forward, 00:01:55/00:02:38

(FC00:2:255::1, FF06:1::1), 00:01:14/00:02:15, flags: SJTgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.6
  Immediate Outgoing interface list:
    Vlan101, Forward, 00:01:14/00:03:18
Leaf-01#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 1:

```
Leaf-01# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
```

```
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF00::/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 412/412/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF00::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF02::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   4/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan101 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:102::12,FF06:1::1) Flags: HW
   SW Forwarding: 3/0/100/0, Other: 0/0/0
   HW Forwarding:   58/0/125/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan101 Flags: F NS
     Pkts: 0/0/3    Rate: 0 pps
 (FC00:2:255::1,FF06:1::1) Flags: HW
   SW Forwarding: 1/0/100/0, Other: 0/0/0
   HW Forwarding:   36/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan101 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,FF10::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF12::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF20::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF22::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF30::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF32::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF33::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF34::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF35::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF36::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF37::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF38::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF39::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3A::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3B::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3C::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3D::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3E::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3F::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF40::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF42::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF50::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF52::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF60::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF62::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF70::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF72::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF80::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF82::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF90::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF92::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFA0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFA2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFB0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFB2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFC0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFC2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
 (*,FFD0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFD2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFE0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFE2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFF0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFF2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
Leaf-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 1:

```
Leaf-01# show bgp ipv6 mvpn all
BGP table version is 43, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                      172.16.255.4             0    100      0 ?
 * i                    172.16.255.4             0    100      0 ?
 *>i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
                      172.16.255.6             0    100      0 ?
 * i                    172.16.255.6             0    100      0 ?
 *>   [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                      ::                                    32768 ?
 *>   [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                      ::                                    32768 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                      ::                                    32768 ?
Leaf-01#
```

The following example shows the output for the **show ip mroute** command on VTEP 1:

```
Leaf-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
```

```
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d01h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:49

(172.16.254.4, 239.1.1.1), 00:01:54/00:01:05, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:01:54/00:01:05

(172.16.254.6, 239.1.1.1), 00:01:56/00:01:03, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:01:56/00:01:03

(*, 224.0.1.40), 1d01h/00:02:53, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d01h/00:02:53

(*, 225.0.0.102), 1d01h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:49

(172.16.254.4, 225.0.0.102), 1d01h/00:02:01, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:10

(172.16.254.6, 225.0.0.102), 1d01h/00:02:20, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:49

(*, 225.0.0.101), 1d01h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:49

(172.16.254.3, 225.0.0.101), 1d01h/00:01:58, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:08
Leaf-01#
```

The following example shows the output for the **show ip mfib** command on VTEP 1:

```
Leaf-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
```

```
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/114/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 13/0/127/0, Other: 2/2/0
   HW Forwarding:   7870/0/164/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   5222/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   2137/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   11/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 4/0/150/0, Other: 0/0/0
   HW Forwarding:   518/0/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
```

```
   HW Forwarding:   498/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
Leaf-01#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 535

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peers** command on VTEP 2:

```
Leaf-02# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/-/4 1d01h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 1d01h
nve1       50901    L3CP 172.16.254.6     0c75.bd67.ef48 50901      UP  A/M/6 1d01h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 1d01h
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  1d01h
nve1       10102    L2CP 172.16.254.6     5              10102      UP   N/A  1d01h
Leaf-02#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
Leaf-02# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ------------------------ ---------- -------- --------
nve1      10101    172.16.254.3             6          10101    1d01h
nve1      10102    172.16.254.6             5          10102    1d01h
Leaf-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on VTEP 2:

```
Leaf-02# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 63, main routing table version 63
6 network entries using 2352 bytes of memory
8 path entries using 1280 bytes of memory
4/4 BGP path/bestpath attribute entries using 1248 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
17 BGP extended community entries using 2372 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7436 total bytes of memory
BGP activity 128/33 prefixes, 221/93 paths, scan interval 60 secs
6 networks peaked at 15:50:41 Aug 6 2020 UTC (21:30:56.871 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1797    1698       63    0    0 1d01h           2
172.16.255.2    4        65001    1792    1701       63    0    0 1d01h           2
Leaf-02#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **group-map** command on VTEP 2:

```
Leaf-02# show ip pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Vl901,::FFFF:172.16.254.6
    Info source: Static
    Uptime: 1d01h, Groups: 1
Leaf-02#
```

The following example shows the output for the **show ip routing vrf** command on VTEP 2:

```
Leaf-02# show ip routing vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "bgp 65001", distance 200, metric 2, type internal
  Route count is 1/1, share count 0
  Routing paths:
    172.16.254.6%default, Vlan901%default
      From AC10:FF01::
      opaque_ptr 0x7F65BA333AD0
      Last updated 1d01h ago
Leaf-02#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on VTEP 2:

```
Leaf-02# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan102
   1d00h     00:02:25
Leaf-02#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 1d00h/never, RP FC00:2:255::255, flags: SCJg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.6
  Immediate Outgoing interface list:
    Vlan102, Forward, 1d00h/never

(FC00:1:102::12, FF06:1::1), 00:05:45/00:01:27, flags: SFJTGq
  Incoming interface: Vlan102
  RPF nbr: FE80::46D3:CAFF:FE28:6CC5
```

```
  Immediate Outgoing interface list:
    Vlan901, Forward, 00:05:45/never

(FC00:2:255::1, FF06:1::1), 00:05:04/00:02:07, flags: SJTgQ
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.6
  Inherited Outgoing interface list:
    Vlan102, Forward, 1d00h/never
Leaf-02#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on VTEP 2:

```
Leaf-02# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF00::/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF00::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF02::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 9/9/0
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   3/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A NS
   Vlan102 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:102::12,FF06:1::1) Flags: HW
   SW Forwarding: 3/0/100/0, Other: 2/2/0
   HW Forwarding:   168/0/118/0, Other: 0/0/0
   Vlan102 Flags: A F
     Pkts: 0/0/0    Rate: 0 pps
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/2    Rate: 0 pps
 (FC00:2:255::1,FF06:1::1) Flags: HW
   SW Forwarding: 1/0/100/0, Other: 0/0/0
   HW Forwarding:   148/0/126/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan102 Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,FF10::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF12::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
(*,FF20::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF22::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF30::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF32::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF33::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF34::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF35::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF36::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF37::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF38::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF39::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3A::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3B::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3C::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3D::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3E::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3F::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF40::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF42::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF50::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF52::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF60::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF62::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
 (*,FF70::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF72::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF80::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF82::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF90::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF92::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFA0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFA2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFB0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFB2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFC0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFC2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFD0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFD2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFE0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFE2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFF0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFF2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
Leaf-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on VTEP 2:

```
Leaf-02# show bgp ipv6 mvpn all
BGP table version is 63, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>   [5][1:1][FC00:1:102::12][FF06:1::1]/42
                       ::                                   32768 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
```

```
                     172.16.255.6             0    100      0 ?
 *>i                  172.16.255.6             0    100      0 ?
 *>   [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                     ::                                    32768 ?
 *>i  [7][1:1][65001][FC00:1:102::12][FF06:1::1]/46
                     172.16.255.3             0    100      0 ?
 *>   [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                     ::                                    32768 ?
Route Distinguisher: 172.16.254.4:102
 * i  [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                     172.16.255.3             0    100      0 ?
 *>i                  172.16.255.3             0    100      0 ?
Leaf-02#
```

The following example shows the output for the **show ip mroute** command on VTEP 2:

```
Leaf-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d01h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:01:32

(172.16.254.4, 239.1.1.1), 00:05:43/00:01:46, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:05:43/00:02:43

(172.16.254.6, 239.1.1.1), 00:05:45/00:01:06, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:05:45/00:00:14

(*, 224.0.1.40), 1d01h/00:02:31, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d01h/00:02:31

(*, 225.0.0.102), 1d01h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:01:32

(172.16.254.6, 225.0.0.102), 1d01h/00:00:55, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
```

```
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:01:32

(172.16.254.4, 225.0.0.102), 1d01h/00:01:49, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:26

(*, 225.0.0.101), 1d01h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:01:32

(172.16.254.3, 225.0.0.101), 1d01h/00:01:46, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:01:32
Leaf-02#
```

The following example shows the output for the **show ip mfib** command on VTEP 2:

```
Leaf-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2/0/170/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 1/0/150/0, Other: 0/0/0
   HW Forwarding:   7870/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
    HW Forwarding:   1/0/224/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,225.0.0.102) Flags: HW
    SW Forwarding: 2/0/163/0, Other: 3/1/2
    HW Forwarding:   5353/0/164/0, Other: 0/0/0
    Null0 Flags: A
    GigabitEthernet1/0/2 Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (172.16.254.6,225.0.0.102) Flags: HW
    SW Forwarding: 1/0/206/0, Other: 0/0/0
    HW Forwarding:   2165/0/163/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
  (*,232.0.0.0/8) Flags: HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   0/0/0/0, Other: 0/0/0
  (*,239.1.1.1) Flags: C HW
    SW Forwarding: 0/0/0/0, Other: 0/0/0
    HW Forwarding:   5/0/168/0, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/0    Rate: 0 pps
  (172.16.254.4,239.1.1.1) Flags: HW
    SW Forwarding: 1/0/150/0, Other: 1495/1491/4
    HW Forwarding:   742/0/156/0, Other: 0/0/0
    Null0 Flags: A NS
    GigabitEthernet1/0/2 Flags: F
      Pkts: 0/0/1    Rate: 0 pps
  (172.16.254.6,239.1.1.1) Flags: HW
    SW Forwarding: 1/0/150/0, Other: 0/0/0
    HW Forwarding:   1460/1/168/1, Other: 0/0/0
    GigabitEthernet1/0/2 Flags: A NS
    Tunnel0, VXLAN Decap Flags: F NS
      Pkts: 0/0/1    Rate: 0 pps
Leaf-02#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 535

Outputs to Verify the Configuration on Border VTEP

The following example shows the output for the **show nve peers** command on Border VTEP:

```
Border# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/-/4 1d01h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/-/4 1d01h
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/6 1d01h
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/6 1d01h
nve1       10101    L2CP 172.16.254.3     6              10101      UP   N/A  1d01h
nve1       10102    L2CP 172.16.254.4     7              10102      UP   N/A  1d00h
Border#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on Border VTEP:

```
Border# show l2vpn evpn peers vxlan
Interface VNI      Peer-IP                  Num routes eVNI     UP time
```

```
--------- -------- ------------------------  ---------- -------- --------
nve1      10101    172.16.254.3              6          10101    1d01h
nve1      10102    172.16.254.4              7          10102    1d00h
Border#
```

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Border VTEP:

```
Border# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.6, local AS number 65001
BGP table version is 62, main routing table version 62
5 network entries using 1960 bytes of memory
8 path entries using 1280 bytes of memory
4/4 BGP path/bestpath attribute entries using 1248 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
1 BGP community entries using 24 bytes of memory
19 BGP extended community entries using 2682 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7354 total bytes of memory
BGP activity 122/28 prefixes, 244/122 paths, scan interval 60 secs
5 networks peaked at 15:42:39 Aug 6 2020 UTC (21:35:36.535 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    1810    1710       62    0    0 1d01h           3
172.16.255.2    4        65001    1810    1704       62    0    0 1d01h           3
Border#
```

The following example shows the output for the **show ip pim vrf** *vrf-name* **group-map** command on Border VTEP:

```
Border# show ip pim vrf green group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Vl2001,FE80::A2B4:39FF:FE21:9183
    Info source: Static
    Uptime: 1d01h, Groups: 1
Border#
```

The following example shows the output for the **show ip routing vrf** *vrf-name* command on Border VTEP:

```
Border# show ip routing vrf green FC00:2:255::255
Routing entry for FC00:2:255::255/128
  Known via "ospf 1", distance 110, metric 2, type intra area
  Redistributing via bgp 65001
  Route count is 1/1, share count 0
  Routing paths:
    FE80::A2B4:39FF:FE21:9183, Vlan2001
      From FE80::A2B4:39FF:FE21:9183
      Last updated 1d01h ago
Border#
```

The following example shows the output for the **show ipv6 mld vrf** *vrf-name* **groups** command on Border VTEP:

```
Border# show ipv6 mld vrf green groups
MLD Connected Group Membership
Group Address                           Interface
    Uptime    Expires
FF06:1::1                               Vlan102
   1d00h     00:04:02
Border#
```

The following example shows the output for the **show ipv6 mroute vrf** *vrf-name* command on Border VTEP:

```
Border# show ipv6 mroute vrf green
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 1d00h/never, RP FC00:2:255::255, flags: SCG
  Incoming interface: Vlan2001
  RPF nbr: FE80::A2B4:39FF:FE21:9183
  Immediate Outgoing interface list:
    Vlan102, Null, 1d00h/never
    Vlan901, Forward, 1d00h/never

(FC00:1:102::12, FF06:1::1), 00:10:24/now, flags: STg
  Incoming interface: Vlan901
  RPF nbr: ::FFFF:172.16.254.4
  Immediate Outgoing interface list:
    Vlan2001, Forward, 00:10:24/00:03:05
  Inherited Outgoing interface list:
    Vlan102, Null, 1d00h/never

(FC00:2:255::1, FF06:1::1), 00:09:43/never, flags: STGq
  Incoming interface: Vlan2001
  RPF nbr: FE80::A2B4:39FF:FE21:9183
  Immediate Outgoing interface list:
    Vlan901, Forward, 00:09:43/never
  Inherited Outgoing interface list:
    Vlan102, Null, 1d00h/never
Border#
```

The following example shows the output for the **show ipv6 mfib vrf** *vrf-name* command on Border VTEP:

```
Border# show ipv6 mfib vrf green
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
```

```
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
VRF green
 (*,FF00::/8) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF00::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF02::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 9/9/0
 (*,FF06:1::1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 3/3/0
   HW Forwarding:   7/0/122/0, Other: 0/0/0
   Vlan2001 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (FC00:1:102::12,FF06:1::1) Flags: HW
   SW Forwarding: 2/0/100/0, Other: 1/0/1
   HW Forwarding:   311/0/125/0, Other: 0/0/0
   Vlan901, VXLAN Decap Flags: A
   Vlan2001 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (FC00:2:255::1,FF06:1::1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   291/0/122/0, Other: 0/0/0
   Vlan2001 Flags: A
   Vlan901, VXLAN v4 Encap (50901, 239.1.1.1) Flags: F
     Pkts: 0/0/0    Rate: 0 pps
 (*,FF10::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF12::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF20::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF22::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF30::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF32::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FF33::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF34::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF35::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF36::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FF37::/32) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
```

```
(*,FF38::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF39::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3A::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3B::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3C::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3D::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3E::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF3F::/32) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF40::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF42::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF50::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF52::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF60::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF62::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF70::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF72::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF80::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF82::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FF90::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FF92::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFA0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFA2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
(*,FFB0::/15) Flags: HW
  SW Forwarding: 0/0/0/0, Other: 0/0/0
  HW Forwarding:   0/0/0/0, Other: 0/0/0
(*,FFB2::/16) Flags:
  SW Forwarding: 0/0/0/0, Other: 0/0/0
```

```
 (*,FFC0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFC2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFD0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFD2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFE0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFE2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
 (*,FFF0::/15) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,FFF2::/16) Flags:
   SW Forwarding: 0/0/0/0, Other: 0/0/0
Border#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Border VTEP:

```
Border# show bgp ipv4 mvpn all
BGP table version is 62, local router ID is 172.16.255.6
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1 (default for vrf green)
 *>i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
 *>   [5][1:1][FC00:2:255::1][FF06:1::1]/42
                      ::                                    32768 ?
 * i  [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                      172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
 * i  [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                      172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>   [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                      ::                                    32768 ?
Border#
```

The following example shows the output for the **show ip mroute** command on Border VTEP:

```
Border# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
```

```
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf,
       e - encap-helper tunnel flag
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 1d01h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:41

(172.16.254.4, 239.1.1.1), 00:10:23/00:02:45, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 00:10:23/00:01:36

(172.16.254.6, 239.1.1.1), 00:10:25/00:03:25, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:10:25/00:02:56

(*, 224.0.1.40), 1d01h/00:02:45, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Loopback0, Forward/Sparse, 1d01h/00:02:45

(*, 225.0.0.102), 1d01h/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:41

(172.16.254.4, 225.0.0.102), 1d01h/00:02:35, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:41

(172.16.254.6, 225.0.0.102), 1d01h/00:03:27, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:02:49, A

(*, 225.0.0.101), 1d01h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:41

(172.16.254.3, 225.0.0.101), 1d01h/00:01:12, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.26.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d01h/00:02:53
Border#
```

The following example shows the output for the **show ip mfib** command on Border VTEP:

```
Border# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
```

```
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client,
                e   - Encap helper tunnel flag.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Loopback0 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   8/0/146/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 7/0/125/0, Other: 0/0/0
   HW Forwarding:   8010/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/7    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1/0/172/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 1/0/154/0, Other: 0/0/0
   HW Forwarding:   5353/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/1    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 11/0/167/0, Other: 7/6/1
   HW Forwarding:   2207/0/151/0, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/10    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   8/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A NS
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
```

```
   SW Forwarding: 4/0/150/0, Other: 0/0/0
   HW Forwarding:   2032/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   Tunnel0, VXLAN Decap Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 2/0/150/0, Other: 4/4/0
   HW Forwarding:   2015/1/156/1, Other: 0/0/0
   Null0 Flags: A
   GigabitEthernet1/0/2 Flags: F
     Pkts: 0/0/1    Rate: 0 pps
Border#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 535

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Spine Switch 1:

```
Spine-01# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 61, main routing table version 61
5 network entries using 1960 bytes of memory
13 path entries using 2080 bytes of memory
3/3 BGP path/bestpath attribute entries using 912 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
16 BGP extended community entries using 2332 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7428 total bytes of memory
BGP activity 257/209 prefixes, 3003/2900 paths, scan interval 60 secs
5 networks peaked at 15:48:28 Aug 6 2020 UTC (21:38:24.468 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    1700    1808       61    0    0 1d01h           3
172.16.255.4    4        65001    1706    1805       61    0    0 1d01h           3
172.16.255.6    4        65001    1713    1813       61    0    0 1d01h           2
Spine-01#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Spine Switch 1:

```
Spine-01# show bgp ipv6 mvpn all
BGP table version is 61, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                     172.16.255.4             0    100      0 ?
 *>i                  172.16.255.4             0    100      0 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
```

```
                     172.16.255.6             0    100      0 ?
 *>i                  172.16.255.6              0    100      0 ?
 *>i  [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                     172.16.255.3             0    100      0 ?
 * i                  172.16.255.3              0    100      0 ?
 * i                  172.16.255.4              0    100      0 ?
 * i  [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                     172.16.255.4             0    100      0 ?
 * i                  172.16.255.3              0    100      0 ?
 *>i                  172.16.255.3              0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                     172.16.255.3             0    100      0 ?
 * i                  172.16.255.3              0    100      0 ?
 * i                  172.16.255.6              0    100      0 ?
Spine-01#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 1:

```
Spine-01# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-01#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 1:

```
Spine-01# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 00:13:12/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.4, 239.1.1.1), 00:11:10/00:01:49, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.14.4
  Outgoing interface list: Null

(172.16.254.6, 239.1.1.1), 00:13:12/00:02:08, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(*, 224.0.1.40), 1w0d/00:02:04, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
```

```
  Outgoing interface list:
    Loopback2, Forward/Sparse, 1w0d/00:02:04

(*, 225.0.0.102), 1w0d/stopped, RP 172.16.255.255, flags: SP
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(172.16.254.6, 225.0.0.102), 00:19:31/00:02:22, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.16.6
  Outgoing interface list: Null

(172.16.254.4, 225.0.0.102), 1d01h/00:01:52, flags: PA
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.14.4
  Outgoing interface list: Null
Spine-01#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 1:

```
Spine-01# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 82/82/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 1/0/206/0, Other: 279/0/279
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/2 Flags: NS
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/4 Flags: NS
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
```

```
   SW Forwarding: 0/0/0/0, Other: 1/0/1
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   1224/0/168/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   GigabitEthernet1/0/4 Flags: A NS
Spine-01#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 535

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp ipv6 mvpn all summary** command on Spine Switch 2:

```
Spine-02# show bgp ipv6 mvpn all summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 61, main routing table version 61
5 network entries using 1960 bytes of memory
13 path entries using 2080 bytes of memory
3/3 BGP path/bestpath attribute entries using 912 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
1 BGP community entries using 24 bytes of memory
16 BGP extended community entries using 2332 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 7428 total bytes of memory
BGP activity 288/240 prefixes, 3108/3005 paths, scan interval 60 secs
5 networks peaked at 15:49:16 Aug 6 2020 UTC (21:40:40.843 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.3    4        65001    1700    1810       61    0    0 1d01h           3
172.16.255.4    4        65001    1711    1803       61    0    0 1d01h           3
172.16.255.5    4        65001       0       0        1    0    0 08:41:01 Idle
172.16.255.6    4        65001    1710    1815       61    0    0 1d01h           2
172.16.255.7    4        65001       0       0        1    0    0 08:40:29 Idle
Spine-02#
```

The following example shows the output for the **show bgp ipv6 mvpn all** command on Spine Switch 2:

```
Spine-02# show bgp ipv6 mvpn all
BGP table version is 61, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 1:1
 * i  [5][1:1][FC00:1:102::12][FF06:1::1]/42
                      172.16.255.4             0    100      0 ?
```

```
 *>i                   172.16.255.4             0    100      0 ?
 * i  [5][1:1][FC00:2:255::1][FF06:1::1]/42
                      172.16.255.6             0    100      0 ?
 *>i                   172.16.255.6             0    100      0 ?
 *>i  [6][1:1][65001][FC00:2:255::255][FF06:1::1]/46
                      172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 * i                   172.16.255.4             0    100      0 ?
 * i  [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
                      172.16.255.4             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 *>i                   172.16.255.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:102
 *>i  [7][172.16.254.4:102][65001][FC00:1:102::12][FF06:1::1]/46
                      172.16.255.3             0    100      0 ?
 * i                   172.16.255.3             0    100      0 ?
 * i                   172.16.255.6             0    100      0 ?
Spine-02#
```

The following example shows the output for the **show ip pim rp mapping** command on Spine Switch 2:

```
Spine-02# show ip pim rp mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
    RP: 172.16.255.255 (?)
Spine-02#
```

The following example shows the output for the **show ip mroute** command on Spine Switch 2:

```
Spine-02# show ip mroute
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
       L - Local, P - Pruned, R - RP-bit set, F - Register flag,
       T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
       X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
       U - URD, I - Received Source Specific Host Report,
       Z - Multicast Tunnel, z - MDT-data group sender,
       Y - Joined MDT-data group, y - Sending to MDT-data group,
       G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
       N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
       Q - Received BGP S-A Route, q - Sent BGP S-A Route,
       V - RD & Vector, v - Vector, p - PIM Joins on route,
       x - VxLAN group, c - PFP-SA cache created entry,
       * - determined by Assert, # - iif-starg configured on rpf intf
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
 Timers: Uptime/Expires
 Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 239.1.1.1), 3d08h/00:03:24, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:24
    GigabitEthernet1/0/1, Forward/Sparse, 1d01h/00:03:06
    GigabitEthernet1/0/4, Forward/Sparse, 1d01h/00:03:02

(172.16.254.4, 239.1.1.1), 00:15:27/00:02:45, flags: T
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 00:15:27/00:03:02
```

```
    GigabitEthernet1/0/1, Forward/Sparse, 00:15:27/00:03:06

(172.16.254.6, 239.1.1.1), 00:15:29/00:02:38, flags: MT
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:15:29/00:03:06
    GigabitEthernet1/0/2, Forward/Sparse, 00:15:29/00:03:24

(*, 224.0.1.40), 1w0d/00:03:27, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:02:31
    GigabitEthernet1/0/1, Forward/Sparse, 1d01h/00:03:27
    GigabitEthernet1/0/4, Forward/Sparse, 1d01h/00:03:08
    Loopback2, Forward/Sparse, 1w0d/00:02:17

(*, 225.0.0.102), 1w0d/00:03:21, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:02
    GigabitEthernet1/0/1, Forward/Sparse, 1d01h/00:02:50
    GigabitEthernet1/0/4, Forward/Sparse, 1d01h/00:03:21

(172.16.254.4, 225.0.0.102), 1d01h/00:01:55, flags: MT
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/4, Forward/Sparse, 1d01h/00:03:21
    GigabitEthernet1/0/1, Forward/Sparse, 1d01h/00:03:02

(172.16.254.6, 225.0.0.102), 1d01h/00:02:03, flags: MT
  Incoming interface: GigabitEthernet1/0/4, RPF nbr 172.16.26.6
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:02
    GigabitEthernet1/0/1, Forward/Sparse, 1d01h/00:03:13

(*, 225.0.0.101), 3d08h/00:03:29, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:05
    GigabitEthernet1/0/1, Forward/Sparse, 1d01h/00:03:07
    GigabitEthernet1/0/4, Forward/Sparse, 1d01h/00:03:29

(172.16.254.3, 225.0.0.101), 1d01h/00:02:39, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:03:05
    GigabitEthernet1/0/4, Forward/Sparse, 1d01h/00:03:29
Spine-02#
```

The following example shows the output for the **show ip mfib** command on Spine Switch 2:

```
Spine-02# show ip mfib
Entry Flags:    C - Directly Connected, S - Signal, IA - Inherit A flag,
                ET - Data Rate Exceeds Threshold, K - Keepalive
                DDE - Data Driven Event, HW - Hardware Installed
                ME - MoFRR ECMP entry, MNE - MoFRR Non-ECMP entry, MP - MFIB
                MoFRR Primary, RP - MRIB MoFRR Primary, P - MoFRR Primary
                MS  - MoFRR  Entry in Sync, MC - MoFRR entry in MoFRR Client.
I/O Item Flags: IC - Internal Copy, NP - Not platform switched,
                NS - Negate Signalling, SP - Signal Present,
                A - Accept, F - Forward, RA - MRIB Accept, RF - MRIB Forward,
                MA - MFIB Accept, A2 - Accept backup,
```

```
                RA2 - MRIB Accept backup, MA2 - MFIB Accept backup

Forwarding Counts: Pkt Count/Pkts per second/Avg Pkt Size/Kbits per second
Other counts:      Total/RPF failed/Other drops
I/O Item Counts:   HW Pkt Count/FS Pkt Count/PS Pkt Count   Egress Rate in pps
Default
 (*,224.0.0.0/4) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,224.0.1.40) Flags: C HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   Loopback2 Flags: F IC NS
     Pkts: 0/0/0    Rate: 0 pps
 (*,225.0.0.101) Flags: C HW
   SW Forwarding: 9/0/112/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/2    Rate: 0 pps
 (172.16.254.3,225.0.0.101) Flags: HW
   SW Forwarding: 4/0/132/0, Other: 0/0/0
   HW Forwarding:   8067/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/1 Flags: A
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/4    Rate: 0 pps
 (*,225.0.0.102) Flags: C HW
   SW Forwarding: 27/0/101/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.4,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   5404/0/176/0, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.6,225.0.0.102) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2214/0/163/0, Other: 0/0/0
   GigabitEthernet1/0/4 Flags: A NS
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
```

```
    Pkts: 0/0/0    Rate: 0 pps
 (*,232.0.0.0/8) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
 (*,239.1.1.1) Flags: C HW
   SW Forwarding: 9/0/150/0, Other: 0/0/0
   HW Forwarding:   0/0/0/0, Other: 0/0/0
   Tunnel1 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/3    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/3    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/3    Rate: 0 pps
 (172.16.254.4,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2629/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/2 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/4 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
 (172.16.254.6,239.1.1.1) Flags: HW
   SW Forwarding: 0/0/0/0, Other: 0/0/0
   HW Forwarding:   2607/1/168/1, Other: 0/0/0
   GigabitEthernet1/0/4 Flags: A
   GigabitEthernet1/0/1 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
   GigabitEthernet1/0/2 Flags: F NS
     Pkts: 0/0/0    Rate: 0 pps
Spine-02#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 535

Outputs to Verify the Configuration on Router 3 (RP Outside the BGP EVPN VXLAN Fabric)

The following example shows the output for the **show ip pim group-map** command on Router 3:

```
R3# show ipv6 pim group-map ff06:1::1
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
    SM, RP: FC00:2:255::255
    RPF: Tu4,FC00:2:255::255 (us)
    Info source: Static
    Uptime: 1d04h, Groups: 1
R3#
```

The following example shows the output for the **show ipv6 mroute** command on Router 3:

```
R3# show ipv6 mroute
Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT, Y - Joined MDT-data group,
       y - Sending to MDT-data group
       g - BGP signal originated, G - BGP Signal received,
       N - BGP Shared-Tree Prune received, n - BGP C-Mroute suppressed,
```

```
       q - BGP Src-Active originated, Q - BGP Src-Active received
       E - Extranet
Timers: Uptime/Expires
Interface state: Interface, State

(*, FF06:1::1), 1d04h/00:03:12, RP FC00:2:255::255, flags: S
  Incoming interface: Tunnel4
  RPF nbr: FC00:2:255::255
  Immediate Outgoing interface list:
    GigabitEthernet0/0/1, Forward, 1d04h/00:03:12
    GigabitEthernet0/0/0, Forward, 1d01h/00:02:45

(FC00:1:102::12, FF06:1::1), 00:18:43/00:03:12, RP FC00:2:255::255, flags: SPR
  Incoming interface: Tunnel4
  RPF nbr: FC00:2:255::255
  Immediate Outgoing interface list:
    GigabitEthernet0/0/1, Null, 00:18:43/00:02:45
    GigabitEthernet0/0/0, Null, 00:18:43/00:03:12

(FC00:1:102::12, FF06:1::1), 00:18:45/00:03:12, flags: S
  Incoming interface: TenGigabitEthernet0/0/0
  RPF nbr: FE80::A2B4:39FF:FE21:9181
  Inherited Outgoing interface list:
    GigabitEthernet0/0/1, Forward, 1d04h/00:03:12

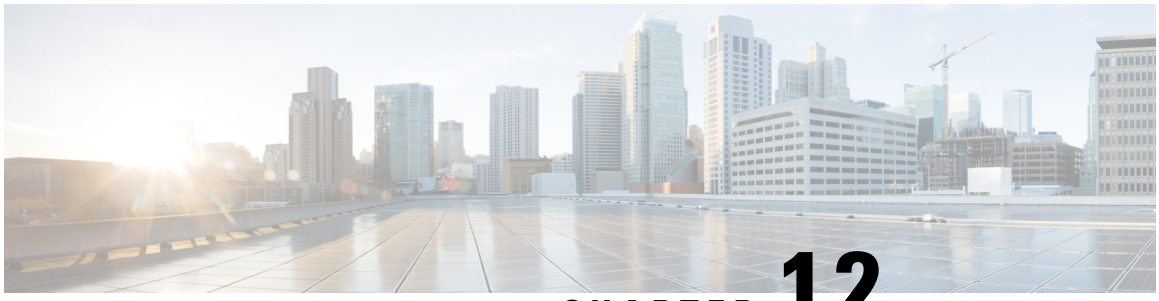
(FC00:2:255::1, FF06:1::1), 00:18:08/00:02:55, RP FC00:2:255::255, flags: SPR
  Incoming interface: Tunnel4
  RPF nbr: FC00:2:255::255
  Immediate Outgoing interface list:
    GigabitEthernet0/0/0, Null, 00:18:08/00:02:45
    GigabitEthernet0/0/1, Null, 00:18:04/00:03:12

(FC00:2:255::1, FF06:1::1), 00:18:06/00:02:55, flags: S
  Incoming interface: GigabitEthernet0/0/0
  RPF nbr: FE80::A2B4:39FF:FE21:9181
  Inherited Outgoing interface list:
   GigabitEthernet0/0/1, Forward, 1d04h/00:03:12
R3#
```

Return to Verifying TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric, on page 535

Troubleshooting Tenant Routed Multicast

See Troubleshoot EVPN VxLAN TRM on Catalyst 9000 Switches document to learn how to troubleshoot issues with TRM in a BGP EVPN VXLAN fabric.



CHAPTER 12

Configuring Multi-Homing in a BGP EVPN VXLAN Fabric

- Restrictions for Multi-Homing in a BGP EVPN VXLAN Fabric, on page 567
- Information About Multi-Homing in a BGP EVPN VXLAN Fabric, on page 567
- How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric, on page 572
- Configuration Examples for Multi-Homing in a BGP EVPN VXLAN Fabric, on page 577

Restrictions for Multi-Homing in a BGP EVPN VXLAN Fabric

- Multi-homing in all-active redundancy mode is not supported.
- Multi-homing in single-active redundancy mode supports only dual-homing, which allows two nodes within a redundancy group.
- Cross-linking between host or access devices and VTEPs is not supported for a dual-homed network.
- A dual-homed network needs internal redundancy to avoid a network split.
- Provision and operational state of EVPN instances must be consistent on both dual-homed VTEPs. Inconsistencies in configuration or operational state of EVPN instances between the VTEPs leads to traffic blackholing.
- A Cisco Industrial Ethernet switches does not support Ethernet Segment configuration if a Leaf node is also a Spine and a BGP Route Reflector.

Information About Multi-Homing in a BGP EVPN VXLAN Fabric

Multi-homing feature in a BGP EVPN VXLAN fabric provides redundancy in the connection between a host or Layer 2 switch and the EVPN VXLAN network.

In a BGP EVPN VXLAN fabric, you connect a host or Layer 2 switch to the EVPN VXLAN network either through single-homing or through multi-homing.

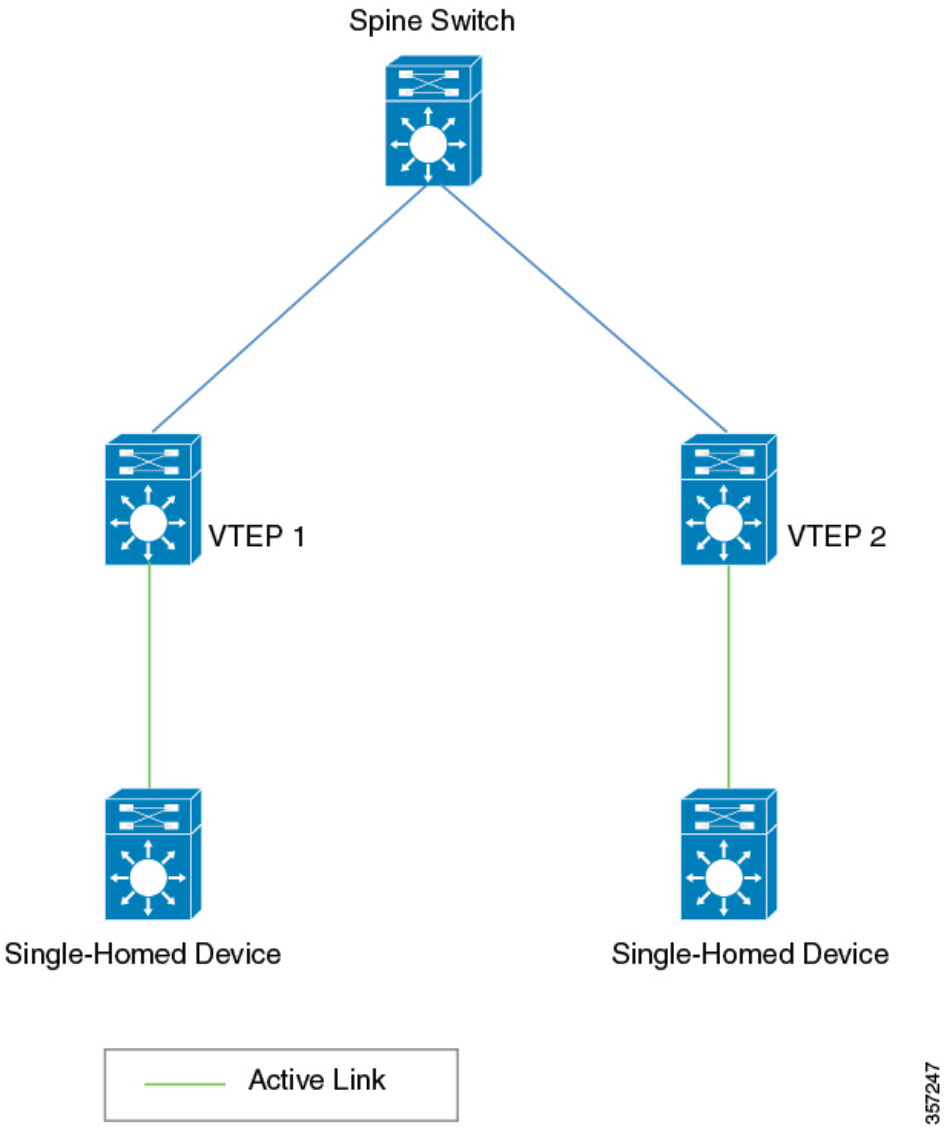
Cisco Industrial Ethernet switches support RFC 7432 and RFC 8365 for VXLAN encapsulation-based EVPN multi-homing capabilities.

Single-Homing

Single-homing allows you to connect a host or Layer 2 switch to a single VTEP in the EVPN VXLAN network. Single-homing does not support redundancy in the connection between the host or access device and the VTEP. When the active link breaks down, the connection between the host (or Layer 2 switch) and the VTEP is lost. As a result, single-homed topologies are not always reliable and efficient.

The following figure shows a single-homed topology:

Figure 45: Single-Homed Topology



Multi-Homing

Multi-homing allows you to connect a host or Layer 2 switch to more than one VTEP in the EVPN VXLAN network. This connection provides redundancy and allows network optimization. Redundancy in the connection with the VTEPs ensures that there's no traffic disruption when there's a network failure. Multi-homed topologies are more reliant, secure, and efficient than single-homed topologies.

Multi-homing operates in single-active and all-active redundancy modes. In both modes, the connected host or access device is represented by an ethernet segment ID. This ethernet segment ID must also be part of the configuration on the VTEP's interface that connects the multihomed host or network device. All traffic forwarded between the VTEPs and the host (or Layer 2 switch) passes through this ethernet segment.

Single-Active Redundancy Mode

In single-active redundancy mode, only a single VTEP, among a group of VTEPs that are attached to a particular Ethernet-segment, is allowed to forward traffic to and from the Ethernet segment. It results in a single-active access link between the VTEPs and the host (or Layer 2 switch) that passes through the Ethernet segment. The single access link can either be a physical link or an ether-channel.

Multi-homing in single-active redundancy mode is supported only in the form of dual-homing. Dual-homing allows a host or access device to be connected to only two VTEPs. A dual-homed topology with single-active redundancy can be deployed in one of the following ways:

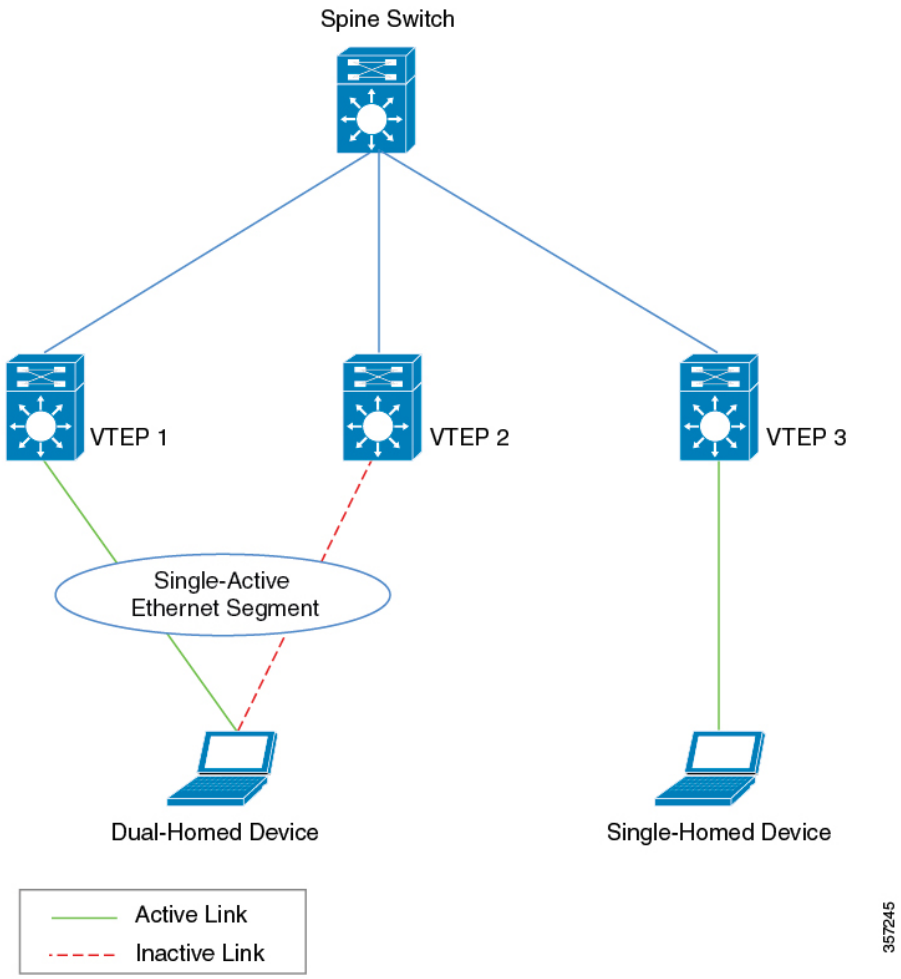
- Dual-homed device
- Dual-homed network

Dual-Homed Device

In single-active dual-homed device topology, a single host or access device is connected to two VTEPs with two links that pass through a single-active Ethernet segment. The Ethernet segment contains two separate links connecting the host or access device to each VTEP, but only one link remains active at any given time. For each VLAN interface on a dual-homed host or access device, only one link remains active. When the active link breaks down, the back-up link takes over and ensures constant connectivity.

The following figure shows a dual-homed device topology:

Figure 46: Dual-Homed Device Topology



Dual-Homed Network

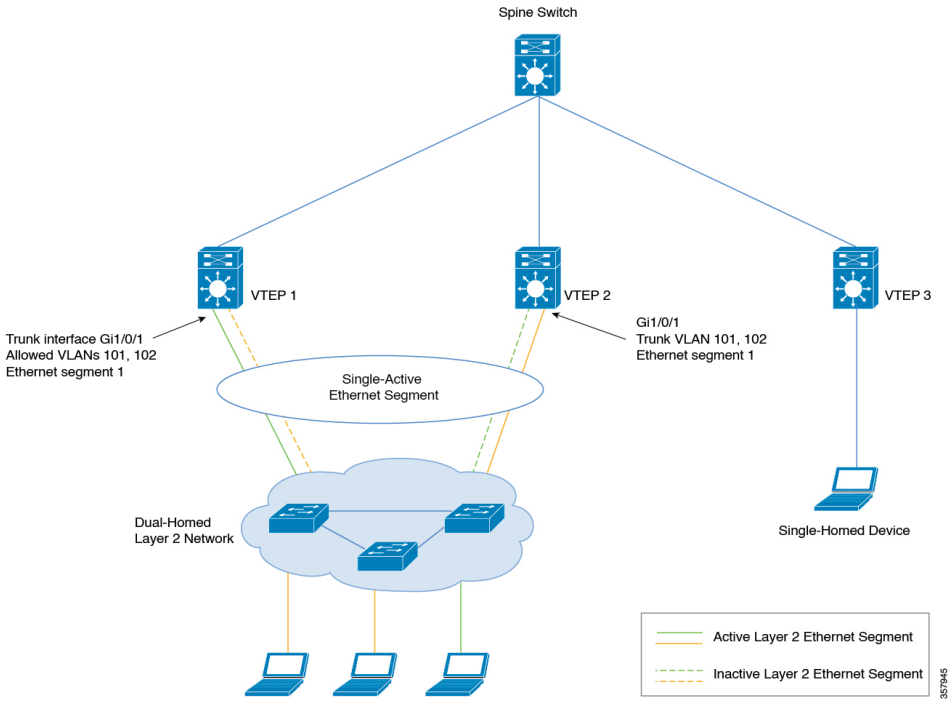
In single-active dual-homed network topology, two host or access devices from the same network are connected to two separate VTEPs through links that pass through a single-active Ethernet segment. At any given time, only one of these links remains active. When the active link breaks down, the back-up link takes over and ensures constant connectivity. The two host or access devices are part of a dual-homed network.

A dual-homed network topology results in a situation where the network splits into two different networks if the connectivity between the host or access devices is lost. To avoid this scenario, redundancy must also be enabled within the dual-homed network.

The following figure shows a dual-homed network topology:

Note Ensure that you enable a spanning tree within the dual-homed Layer 2 network.

Figure 47: Dual-Homed Network Topology



DF Election and Load Balance

A dual-homed network with a single-active Ethernet segment uses a Designated Forwarder (DF) election mechanism to load balance the traffic. The DF election is made at the Layer 2 VNI level, when the access interface from the VTEP is a trunk interface and an Ethernet segment is configured.

In the above topology, some Layer 2 VNIs use the interface connected to VTEP 1 as the active link and the others use the interface connected to VTEP 2 as the active link. This allows effective utilization of bandwidth on both the interfaces in a steady network state. Traffic in each Layer 2 VNI is load balanced for the downstream dual-homed Layer 2 network. If any of the physical interface link to the downstream Layer 2 device goes down and is not operational, the DF election algorithm recalculates the active link interface. After the link is reestablished and both links are operational again, the DF election algorithm restores the load balancing operation to utilize the bandwidth of both the links effectively.

To support selective migration of traditional Layer 2 domains to EVPN overlay Layer 2 segments, the EVPN-enabled VLANs and non EVPN-enabled VLANs can continue to exist on the ethernet segment-enabled trunk interfaces. The DF or non DF forwarding state of the physical port is applicable only for the EVPN-enabled VLANS and will be managed by the EVPN control plane. Spanning Tree Protocol (STP) if enabled, will manage the forwarding state of the physical port only for non EVPN-enabled VLANs.

The Layer 2 trunk port disables classic STP function for the VLANs mapped to EVI once the trunk port is configured with Ethernet Segment (ES). STP functionality will be disabled on per VP (vlan,port) as per the loop-prevention technique and handled by BGP EVPN control-plane with DF/NDF selection algorithm. Such architecture change remains unknown at downstream Layer 2 access network device(s) which continue to operate with STP protocol.

As the data switching for each VLAN is evenly load-share across each Layer 2 trunk port, the fault detection and recovery mechanics between Distribution layer is handled by BGP EVPN control-plane. The downstream

network convergence from IP core to Layer 2 network is rapidly detected, synchronized, and it delivers deterministic sub-second recovery. However, the Upstream network from Layer 2 network to an IP core demands re-learning MAC move from failed connection to new one, that by default L2 Access switch relies on slow STP forward-delay timer.

To deliver bi-directional and symmetric data convergence during Layer 2 path failure you can conditionally inject STP Topology Change Notification (TCN) message to downstream Layer 2 network devices. This event-driven STP TCN forces the Layer 2 network devices to perform rapid MAC move from failed to new connection and start performing data switching to deliver deterministic sub-second convergence time.

Migration Between Single-Homed and Multi-Homed Network Topologies

BGP EVPN VXLAN allows you to migrate your network topology from one redundancy mode to another. You can transition from a single-homed topology to a multi-homed topology. Likewise, you can also remove the redundancy from a multi-homed topology to move back to a single-homed topology.

Note When you migrate from one topology to another, ensure you make corresponding changes to the Ethernet segment configuration. If you change either of the two without making corresponding changes to the other, it results in traffic loops and traffic blackholing.

For a sample illustration and detailed steps about how to migrate from a single-homed topology to a single-active dual-homed topology, see Migrating from a Single-Homed Topology to a Single-Active Dual-Homed Topology, on page 574.

For a sample illustration and detailed steps about how to migrate from a single-active dual-homed topology to a single-homed topology, see Migrating from a Single-Active Dual-Homed Topology to a Single-Homed Topology, on page 576.

How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric

Before you configure multi-homing in a BGP EVPN VXLAN fabric, ensure that you configure EVPN VXLAN Layer 2 and Layer 3 overlay networks. See How to Configure EVPN VXLAN Integrated Routing and Bridging, on page 108 for detailed steps.

Configuring Dual-Homing with Single Active Redundancy in a BGP EVPN VXLAN Fabric

To configure dual-homing with single-active redundancy in a BGP EVPN VXLAN fabric, perform the following set of procedures:

Configuring Redundancy on an Ethernet Segment

To configure redundancy on an ethernet segment, perform the following steps :

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**l2vpn evpn ethernet-segment** *ethernet-segment-id* **Example:** `Device(config)# l2vpn evpn ethernet-segment 1`	Enters Layer 2 VPN EVPN ethernet segment configuration mode.
Step 4	**identifier type** {**0** *esi-value* | **3 system-mac** *mac-addrress*} **Example:** `Device(config-evpn-es)# identifier type 0 0.0.0.0.0.0.0.0.1`	Configures the ethernet segment identifier type (ESI) and value for the ethernet segment. The following ESI types are supported: • Type 0: This type indicates an arbitrary 9-octet ESI value. The format is 00 + 9-octets of ESI value. • Type 3: This type indicates a MAC-based ESI Value. The format is 03 + system-mac (6 bytes) + value of MAC address (3 bytes).
Step 5	**redundancy** *redundancy-type* **Example:** `Device(config-evpn-es)# redundancy single-active`	Configures the redundancy type for the ethernet segment.
Step 6	**df-election wait-time** *time-period* **Example:** `Device(config-evpn-es)# df-election wait-time 1`	(Optional) Configures the designated forwarder (DF) election wait time for the ethernet segment. The range is 1 to 10 seconds. The default wait time is 3 seconds.
Step 7	**end** **Example:** `Device(config-evpn-es)# end`	Exits Layer 2 VPN EVPN ethernet segment configuration mode and enters privileged EXEC mode.

Associating an Ethernet Segment with an Interface on a VTEP

To associate the ethernet segment with an interface on a VTEP, perform the following steps :

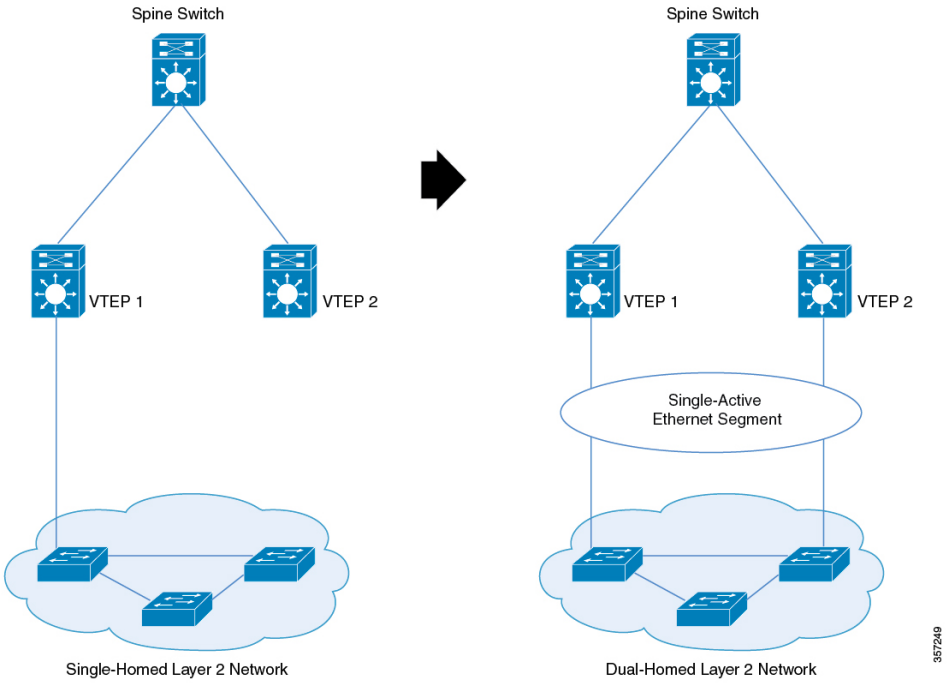
Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *interface-id* **Example:** `Device(config)# interface GigabitEthernet1/0/1`	Specifies the interface, and enters interface configuration mode.
Step 4	**evpn ethernet-segment** *ethernet-segment-id* **Example:** `Device(config-if)# evpn ethernet-segment 1`	Associates the specified Ethernet segment with the interface. Each Ethernet segment is represented by a unique Ethernet segment ID. **Note** Ensure that you configure a unique Ethernet segment ID on any interface. Ensure that you configure the same segment ID on the link that connects the second VTEP and the dual-homed device (the second link through the Ethernet segment).
Step 5	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and enters privileged EXEC mode.

Migrating from a Single-Homed Topology to a Single-Active Dual-Homed Topology

The following figures illustrate the migration from a single-homed topology to a single-active dual-homed topology:

Figure 48: Migration from a Single-Homed Network to a Single-Active Dual-Homed Network



To migrate from a single-homed network to a single-active dual-homed network, perform the following steps:

Note When you migrate from one topology to another, ensure you make corresponding changes to the Ethernet segment configuration. If you change either of the two without making corresponding changes to the other, it results in traffic loops and traffic blackholing.

1. Before you migrate, we recommend that you do not configure the VTEP as the root bridge of the spanning tree, as the provision of Ethernet segment on the interface of a VTEP excludes it from spanning-tree. If the VTEP is the root bridge, its exclusion from the spanning-tree triggers an immediate spanning tree re-convergence.

Note Do not activate a link between VTEP 2 and a switch in the single-homed network yet. Activate the second only once you configure the Ethernet segment. In case a link is already activated, ensure that you deactivate the link.

2. Provision the Ethernet segment on the interface of the VTEP that has the active link. Provision of the Ethernet segment updates all the MAC addresses that are locally learned on that interface with the Ethernet segment ID of the interface.
3. Now provision the Ethernet segment on the interface of the second VTEP that needs to be connected to the switch in the single-homed network.

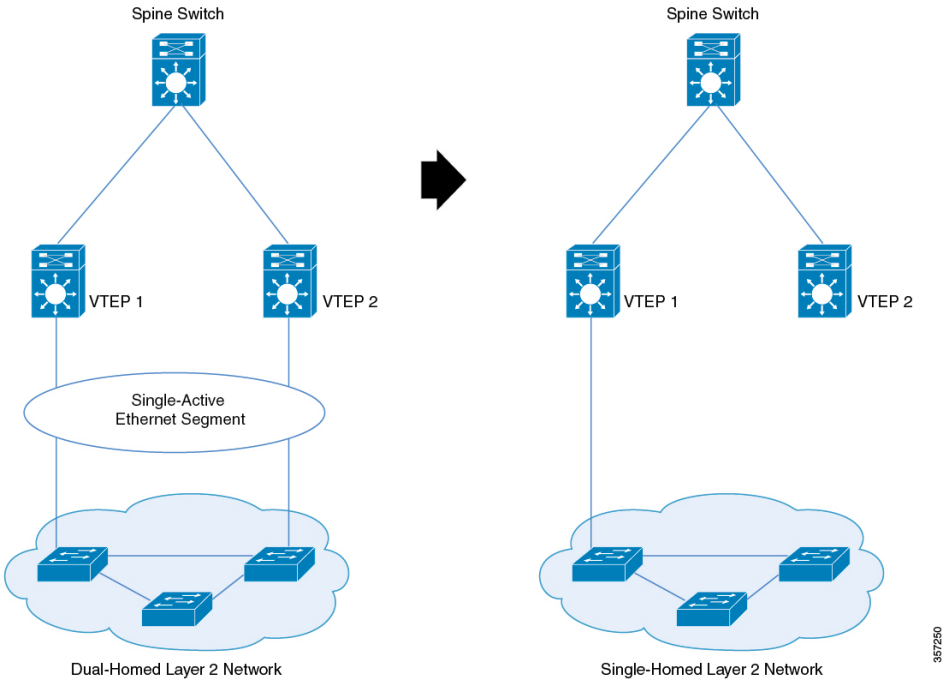
4. Connect the link and bring up the interface of the second VTEP. By doing this, you trigger fast convergence, Ethernet segment auto-discovery, and DF reelection.

   The single-homed network has now migrated to a dual-homed network.

Migrating from a Single-Active Dual-Homed Topology to a Single-Homed Topology

The following figures illustrate the migration from a single-active dual-homed topology to a single-homed topology:

Figure 49: Migration from a Single-Active Dual-Homed Network to a Single-Homed Network



To migrate from a single-active dual-homed network to a single-homed network, perform the following steps:

Note When you migrate from one topology to another, ensure you make corresponding changes to the Ethernet segment configuration. If you change either of the two without making corresponding changes to the other, it results in traffic loops and traffic blackholing.

Note Ensure that the Ethernet segment remains configured on the dual-homed links as long as the links are up. If the Ethernet segment is removed from an active link, it causes traffic loops.

1. Before you migrate, we recommended that you configure portfast on the link you activate. Removal of Ethernet segment from an interface on the VTEP puts it back into the spanning-tree. If the interface is not configured with portfast, the port goes through block-learn-forward states and causes extensive traffic loss.
2. Shut down the interface that needs to be decommissioned. When you shut down the interface, you trigger fast convergence, Ethernet segment auto-discovery, and DF reelection. As a result, all the traffic converges into the active link.
3. Remove the Ethernet segment from the decommissioned interface. Ensure that the interface is down before you disconnect the link.

   The dual-homed network has now migrated to a single-homed network with an Ethernet segment.
4. (Optional) Remove the Ethernet segment from the interface with the active link on the VTEP.

   Removal of the Ethernet segment updates all the MAC addresses that are locally learned on that interface without the Ethernet segment ID.

   The Ethernet segment is now removed from the single-homed network homed network.

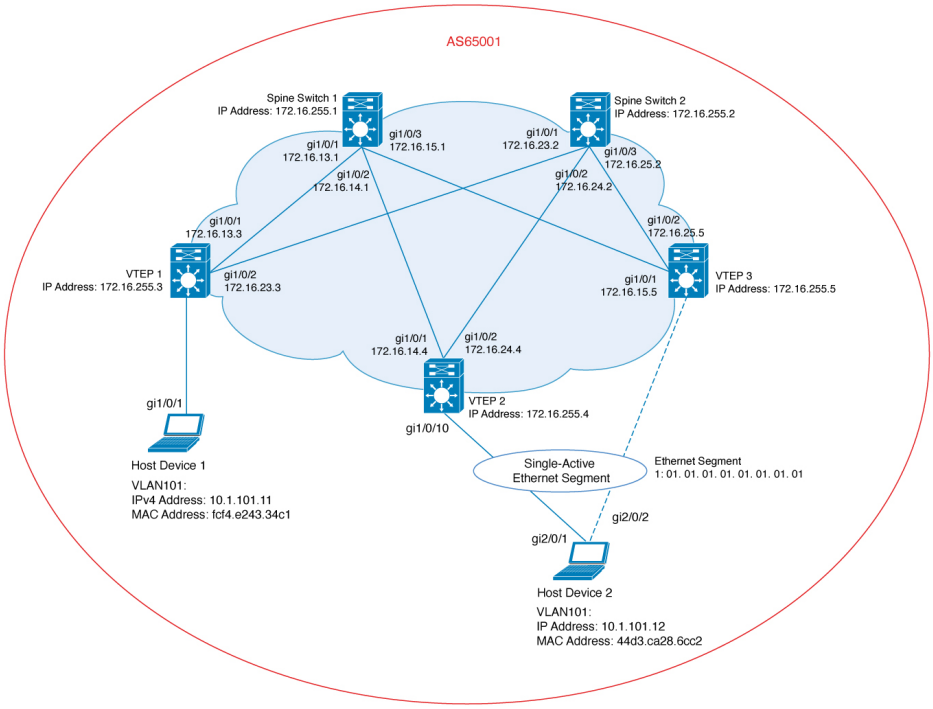
Configuration Examples for Multi-Homing in a BGP EVPN VXLAN Fabric

This section provides configuration examples for multi-homing in a BGP EVPN VXLAN fabric:

Example: Configuring Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric

This example shows how to configure and verify dual-homing with single-active redundancy in a BGP EVPN VXLAN fabric for the following topology:

Figure 50: Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric



The topology shows an EVPN VXLAN network with two spine switches (Spine Switch 1 and Spine Switch 2) and three VTEPs (VTEP 1, VTEP 2, and VTEP 3). Host Device 1 is connected to VTEP 1. Host Device 2 is connected to VTEP 2 and VTEP 3 as a dual-homed single-active connection that passes through Ethernet Segment 1.

Note Ensure that you configure a unique Ethernet segment ID on any interface in the fabric. If an Ethernet segment ID is associated with one of the connecting links passing through the segment, associate the same Ethernet segment ID with the second link.

Note Do not configure a unique Ethernet segment ID per EVPN instance or VLAN or virtual network instance (VNI). For example purpose, EVPN instance 101 is used in the Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric, on page 583 section.

Table 53: Configuring Dual-Homing with Single-Active Redundancy using VTEP 2 and VTEP 3

VTEP 2	VTEP 3
Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn ethernet-segment 1 identifier type 0 01.01.01.01.01.01.01.01.01 redundancy single-active ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 !	LEaf-03# **show running-config** hostname Leaf-03 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn ethernet-segment 1 identifier type 0 01.01.01.01.01.01.01.01.01 redundancy single-active ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.5 255.255.255.255 ip ospf 1 area 0 !

VTEP 2	VTEP 3
interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access evpn ethernet-segment 1 spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 no autostate ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 50901 vrf green	interface Loopback1 ip address 172.16.254.5 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.15.5 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.25.5 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 switchport access vlan 101 switchport mode access evpn ethernet-segment 1 spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 50901 vrf green

VTEP 2	VTEP 3
!	!
router ospf 1	router ospf 1
router-id 172.16.255.4	router-id 172.16.255.5
!	!
router bgp 65001	router bgp 65001
bgp log-neighbor-changes	bgp log-neighbor-changes
no bgp default ipv4-unicast	no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001	neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source Loopback0	neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.2 remote-as 65001	neighbor 172.16.255.2 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0	neighbor 172.16.255.2 update-source Loopback0
!	!
address-family ipv4	address-family ipv4
exit-address-family	exit-address-family
!	!
address-family l2vpn evpn	address-family l2vpn evpn
neighbor 172.16.255.1 activate	neighbor 172.16.255.1 activate
neighbor 172.16.255.1 send-community both	neighbor 172.16.255.1 send-community both
neighbor 172.16.255.2 activate	neighbor 172.16.255.2 activate
neighbor 172.16.255.2 send-community both	neighbor 172.16.255.2 send-community both
exit-address-family	exit-address-family
!	!
address-family ipv4 vrf green	address-family ipv4 vrf green
advertise l2vpn evpn	advertise l2vpn evpn
redistribute connected	redistribute connected
redistribute static	redistribute static
exit-address-family	exit-address-family
!	!
end	end
!	!
Leaf-02#	Leaf-03#

Table 54: Configuring Spine Switch 1, Spine Switch 2, and VTEP 1 to Configure Dual-Homing with Single-Active Redundancy

Spine Switch 1	Spine Switch 2	VTEP 1
Spine-01# **show running-config** hostname Spine-01 ! ip routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 no switchport ip address 172.16.15.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.1 !	Spine-01# **show running-config** hostname Spine-01 ! ip routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 no switchport ip address 172.16.15.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! router ospf 1 router-id 172.16.255.1 !	Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip ospf 1 area 0 !

Spine Switch 1	Spine Switch 2	VTEP 1
router bgp 65001 bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 neighbor 172.16.255.5 remote-as 65001 neighbor 172.16.255.5 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.5 activate neighbor 172.16.255.5 send-community both neighbor 172.16.255.5 route-reflector-client exit-address-family ! end ! Spine-01#	router bgp 65001 bgp router-id 172.16.255.2 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.3 remote-as 65001 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65001 neighbor 172.16.255.4 update-source Loopback0 neighbor 172.16.255.5 remote-as 65001 neighbor 172.16.255.5 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.5 activate neighbor 172.16.255.5 send-community both neighbor 172.16.255.5 route-reflector-client exit-address-family ! end ! Spine-02#	interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 ingress-replication member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.3 !

Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric

The following sections provide sample outputs for **show** commands to verify dual-homing with single-active redundancy on the devices in the topology configured above:

- Outputs to Verify the Configuration on VTEP 1, on page 584
- Outputs to Verify the Configuration on VTEP 2, on page 586
- Outputs to Verify the Configuration on VTEP 3, on page 589

• Outputs to Verify the Configuration on Spine Switch 1, on page 593
• Outputs to Verify the Configuration on Spine Switch 2, on page 595

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show nve peer** command on VTEP 1:

```
Leaf-01# show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.5     7c21.0dbd.2748 50901      UP  A/M/4 01:17:04
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/4 03:26:09
nve1       10101    L2CP 172.16.254.4     8              10101      UP   N/A  03:52:15
nve1       10101    L2CP 172.16.254.5     10             10101      UP   N/A  05:25:28

Leaf-01#
```

The following example shows the output for the **show l2vpn evpn evi** *evpn-instance* **detail** command on VTEP 1:

```
Leaf-01# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.3:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Vlan:              101
    Ethernet-Tag:    0
    State:           Established
    Core If:         Vlan901
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            10b3.d56a.8fc8
    Core Vlan:       901
    L2 VNI:          10101
    L3 VNI:          50901
    VTEP IP:         172.16.254.3
    VRF:             green
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.4
        Routes: 4 MAC, 2 MAC/IP, 1 IMET, 1 EAD
      172.16.254.5
        Routes: 6 MAC, 2 MAC/IP, 1 IMET, 1 EAD

Leaf-01#
```

The following example shows the output for the **show bgp l2vpn evpn evi** *evpn-instance* command on VTEP 1:

```
Leaf-01# show bgp l2vpn evpn evi 101
BGP table version is 6958, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>i  [1][172.16.254.3:101][00010101010101010101][0]/23
                      172.16.254.5             0    100      0 ?
 *mi                   172.16.254.4             0    100      0 ?
 *>   [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      ::                                    32768 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286C82][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      ::                                    32768 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      ::                                    32768 ?
 *>   [3][172.16.254.3:101][0][32][172.16.254.3]/17
                      ::                                    32768 ?
 *>i  [3][172.16.254.3:101][0][32][172.16.254.4]/17
                      172.16.254.4             0    100      0 ?
 *>i  [3][172.16.254.3:101][0][32][172.16.254.5]/17
                      172.16.254.5             0    100      0 ?

Leaf-01#
```

The following example shows the output for the **show l2route evpn mac** command on VTEP 1:

```
Leaf-01# show l2route evpn mac
  EVI       ETag  Prod    Mac Address                    Next Hop(s) Seq Number
----- ---------- ----- -------------- ------------------------------ ----------
  101          0 L2VPN 10b3.d56a.8fc1                       Vl101:0          0
  101          0   BGP 44d3.ca28.6c82         V:10101 172.16.254.5          0
  101          0   BGP 44d3.ca28.6cc2         V:10101 172.16.254.5          0
  101          0   BGP 7c21.0dbd.2741         V:10101 172.16.254.5          0
  101          0   BGP 7c21.0dbd.9541         V:10101 172.16.254.4          0
  101          0 L2VPN f4cf.e243.34c1                 Gi1/0/10          0

Leaf-01#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* command on VTEP 1:

```
Leaf-01# show l2route evpn mac esi 0001.0101.0101.0101.0101
  EVI       ETag  Prod    Mac Address                    Next Hop(s) Seq Number
----- ---------- ----- -------------- ------------------------------ ----------
  101          0   BGP 44d3.ca28.6c82         V:10101 172.16.254.5          0
  101          0   BGP 44d3.ca28.6cc2         V:10101 172.16.254.5          0

Leaf-01#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* **detail** command on VTEP 1:

```
Leaf-01# show l2route evpn mac esi 0001.0101.0101.0101.0101 detail
EVPN Instance:            101
Ethernet Tag:             0
Producer Name:            BGP
MAC Address:              44d3.ca28.6c82
Num of MAC IP Route(s):   0
Sequence Number:          0
ESI:                      0001.0101.0101.0101.0101
Flags:                    B()
Next Hop(s):              V:10101 172.16.254.5
Resolved Next Hops:       V:10101 172.16.254.5, V:10101 172.16.254.4
Resolved Redundancy Mode: Single-Active

EVPN Instance:            101
Ethernet Tag:             0
Producer Name:            BGP
MAC Address:              44d3.ca28.6cc2
Num of MAC IP Route(s):   0
Sequence Number:          0
ESI:                      0001.0101.0101.0101.0101
Flags:                    B()
Next Hop(s):              V:10101 172.16.254.5
Resolved Next Hops:       V:10101 172.16.254.5, V:10101 172.16.254.4
Resolved Redundancy Mode: Single-Active

Leaf-01#
```

Return to Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric, on page 583.

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show nve peer** command on VTEP 2:

```
Leaf-02# show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/4 03:24:45
nve1       50901    L3CP 172.16.254.5     7c21.0dbd.2748 50901      UP  A/M/4 01:15:39
nve1       10101    L2CP 172.16.254.3     5              10101      UP   N/A  03:24:45
nve1       10101    L2CP 172.16.254.5     6              10101      UP   N/A  03:24:45

Leaf-02#
```

The following example shows the output for the **show l2vpn evpn ethernet-segment detail** command on VTEP 2:

```
Leaf-02# show l2vpn evpn ethernet-segment detail
EVPN Ethernet Segment ID: 0001.0101.0101.0101.0101
  Interface:              Gi1/0/10
  Redundancy mode:        single-active
  DF election wait time:  3 seconds
  Split Horizon label:    0
  State:                  Ready
  Encapsulation:          vxlan
  Ordinal:                0
  RD:                     172.16.254.4:7
    Export-RTs:           65001:101
```

```
  Forwarder List:         172.16.254.4 172.16.254.5

Leaf-02#
```

The following example shows the output for the **show l2vpn evpn evi** *evpn-instance* **detail** command on VTEP 2:

```
Leaf-02# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Vlan:              101
    Ethernet-Tag:    0
    State:           Established
    Core If:         Vlan901
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            7c21.0dbd.9548
    Core Vlan:       901
    L2 VNI:          10101
    L3 VNI:          50901
    VTEP IP:         172.16.254.4
    VRF:             green
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101 (DF state: blocked)
        Routes: 0 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.5
        Routes: 3 MAC, 1 MAC/IP, 1 IMET, 1 EAD

Leaf-02#
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 4 on VTEP 2:

```
Leaf-02# show bgp l2vpn evpn route-type 4
BGP routing table entry for [4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23,
 version 601
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      Local vtep: 172.16.254.4
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:41:40 UTC
BGP routing table entry for [4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23,
 version 658
```

```
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 6
  Local
    172.16.254.5 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.5, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Jan 26 2021 19:43:19 UTC
  Refresh Epoch 6
  Local
    172.16.254.5 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.5, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:43:19 UTC

Leaf-02#
```

The following example shows the output for the **show bgp l2vpn evpn evi** *evpn-instance* command on VTEP 2:

```
Leaf-02# show bgp l2vpn evpn evi 101
BGP table version is 845, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.4:101
 *mi  [1][172.16.254.4:101][00010101010101010101][0]/23
                       172.16.254.5             0    100      0 ?
 *>                     ::                                    32768 ?
 *>i  [2][172.16.254.4:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                       172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286C82][0][*]/20
                       172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                       172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                       172.16.254.5             0    100      0 ?
 *>   [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                       ::                                    32768 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][0][*]/20
                       172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                       172.16.254.3             0    100      0 ?
 *>i  [3][172.16.254.4:101][0][32][172.16.254.3]/17
                       172.16.254.3             0    100      0 ?
 *>   [3][172.16.254.4:101][0][32][172.16.254.4]/17
                       ::                                    32768 ?
 *>i  [3][172.16.254.4:101][0][32][172.16.254.5]/17
                       172.16.254.5             0    100      0 ?

Leaf-02#
```

The following example shows the output for the **show l2route evpn mac** command on VTEP 2:

```
Leaf-02# show l2route evpn mac
  EVI       ETag  Prod    Mac Address                    Next Hop(s) Seq Number
----- ---------- ----- -------------- ------------------------------ ----------
  101          0   BGP 10b3.d56a.8fc1         V:10101 172.16.254.3          0
  101          0   BGP 44d3.ca28.6c82         V:10101 172.16.254.5          0
  101          0   BGP 44d3.ca28.6cc2         V:10101 172.16.254.5          0
  101          0   BGP 7c21.0dbd.2741         V:10101 172.16.254.5          0
  101          0 L2VPN 7c21.0dbd.9541                      Vl101:0          0
  101          0   BGP f4cf.e243.34c1         V:10101 172.16.254.3          0

Leaf-02#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* command on VTEP 2:

```
Leaf-02# show l2route evpn mac esi 0001.0101.0101.0101.0101
  EVI       ETag  Prod    Mac Address                    Next Hop(s) Seq Number
----- ---------- ----- -------------- ------------------------------ ----------
  101          0   BGP 44d3.ca28.6c82         V:10101 172.16.254.5          0
  101          0   BGP 44d3.ca28.6cc2         V:10101 172.16.254.5          0

Leaf-02#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* **detail** command on VTEP 2:

```
Leaf-02# show l2route evpn mac esi 0001.0101.0101.0101.0101 detail
EVPN Instance:            101
Ethernet Tag:             0
Producer Name:            BGP
MAC Address:              44d3.ca28.6c82
Num of MAC IP Route(s):   0
Sequence Number:          0
ESI:                      0001.0101.0101.0101.0101
Flags:                    B()
Next Hop(s):              V:10101 172.16.254.5
Resolved Next Hops:       V:10101 172.16.254.5
Resolved Redundancy Mode: Single-Active

EVPN Instance:            101
Ethernet Tag:             0
Producer Name:            BGP
MAC Address:              44d3.ca28.6cc2
Num of MAC IP Route(s):   0
Sequence Number:          0
ESI:                      0001.0101.0101.0101.0101
Flags:                    B()
Next Hop(s):              V:10101 172.16.254.5
Resolved Next Hops:       V:10101 172.16.254.5
Resolved Redundancy Mode: Single-Active

Leaf-02#
```

Return to Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric, on page 583.

Outputs to Verify the Configuration on VTEP 3

The following example shows the output for the **show nve peer** command on VTEP 3:

```
Leaf-03# show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/4 04:23:46
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/4 03:24:57
nve1       10101    L2CP 172.16.254.3     5              10101      UP   N/A  04:23:46
nve1       10101    L2CP 172.16.254.4     4              10101      UP   N/A  03:24:57

Leaf-03#
```

The following example shows the output for the **show l2vpn evpn ethernet-segment detail** command on VTEP 3:

```
Leaf-03# show l2vpn evpn ethernet-segment detail
EVPN Ethernet Segment ID: 0001.0101.0101.0101.0101
  Interface:              Gi1/0/10
  Redundancy mode:        single-active
  DF election wait time:  3 seconds
  Split Horizon label:    0
  State:                  Ready
  Encapsulation:          vxlan
  Ordinal:                1
  RD:                     172.16.254.5:9
    Export-RTs:           65001:101
  Forwarder List:         172.16.254.4 172.16.254.5

Leaf-03#
```

The following example shows the output for the **show l2vpn evpn evi** *evpn-instance* **detail** command on VTEP 3:

```
Leaf-03# show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.5:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Vlan:              101
    Ethernet-Tag:    0
    State:           Established
    Core If:         Vlan901
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            7c21.0dbd.2748
    Core Vlan:       901
    L2 VNI:          10101
    L3 VNI:          50901
    VTEP IP:         172.16.254.5
    VRF:             green
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101 (DF state: forwarding)
        Routes: 2 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.4
```

```
        Routes: 1 MAC, 1 MAC/IP, 1 IMET, 1 EAD

Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 4 on VTEP 3:

```
Leaf-03# show bgp l2vpn evpn route-type 4
BGP routing table entry for [4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23,
 version 337
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 5
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Jan 26 2021 19:38:35 UTC
  Refresh Epoch 5
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:38:35 UTC
BGP routing table entry for [4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23,
 version 1269
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.5)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      Local vtep: 172.16.254.5
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:40:14 UTC

Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn evi** *evpn-instance* command on VTEP 3:

```
Leaf-03# show bgp l2vpn evpn evi 101
BGP table version is 1284, local router ID is 172.16.255.5
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.5:101
 *>   [1][172.16.254.5:101][00010101010101010101][0]/23
                       ::                                    32768 ?
 *mi                    172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.5:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
```

```
                    172.16.254.3             0    100      0 ?
 *>   [2][172.16.254.5:101][0][48][44D3CA286C82][0][*]/20
                    ::                                    32768 ?
 *>   [2][172.16.254.5:101][0][48][44D3CA286CC2][0][*]/20
                    ::                                    32768 ?
 *>   [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                    ::                                    32768 ?
 *>i  [2][172.16.254.5:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                    172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.5:101][0][48][F4CFE24334C1][0][*]/20
                    172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.5:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                    172.16.254.3             0    100      0 ?
 *>i  [3][172.16.254.5:101][0][32][172.16.254.3]/17
                    172.16.254.3             0    100      0 ?
 *>i  [3][172.16.254.5:101][0][32][172.16.254.4]/17
                    172.16.254.4             0    100      0 ?
 *>   [3][172.16.254.5:101][0][32][172.16.254.5]/17
                    ::                                    32768 ?
Leaf-03#
```

The following example shows the output for the **show l2route evpn mac** command on VTEP 3:

```
Leaf-03# show l2route evpn mac
  EVI       ETag  Prod    Mac Address                  Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------
  101          0   BGP 10b3.d56a.8fc1         V:10101 172.16.254.3          0
  101          0 L2VPN 44d3.ca28.6c82                 Gi1/0/10          0
  101          0 L2VPN 44d3.ca28.6cc2                 Gi1/0/10          0
  101          0 L2VPN 7c21.0dbd.2741                      Vl101:0          0
  101          0   BGP 7c21.0dbd.9541         V:10101 172.16.254.4          0
  101          0   BGP f4cf.e243.34c1         V:10101 172.16.254.3          0

Leaf-03#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* command on VTEP 3:

```
Leaf-03# show l2route evpn mac esi 0001.0101.0101.0101.0101
  EVI       ETag  Prod    Mac Address                  Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------
  101          0 L2VPN 44d3.ca28.6c82                 Gi1/0/10          0
  101          0 L2VPN 44d3.ca28.6cc2                 Gi1/0/10          0

Leaf-03#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* **detail** command on VTEP 3:

```
Leaf-03# show l2route evpn mac esi 0001.0101.0101.0101.0101 detail
EVPN Instance:            101
Ethernet Tag:             0
Producer Name:            L2VPN
MAC Address:              44d3.ca28.6c82
Num of MAC IP Route(s):   0
Sequence Number:          0
ESI:                      0001.0101.0101.0101.0101
Flags:                    B()
Next Hop(s):              Gi1/0/10

EVPN Instance:            101
```

```
Ethernet Tag:              0
Producer Name:             L2VPN
MAC Address:               44d3.ca28.6cc2
Num of MAC IP Route(s):    0
Sequence Number:           0
ESI:                       0001.0101.0101.0101.0101
Flags:                     B()
Next Hop(s):               Gi1/0/10

Leaf-03#
```

Return to Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric, on page 583.

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 5443, main routing table version 5443
17 network entries using 5848 bytes of memory
34 path entries using 7072 bytes of memory
13/11 BGP path/bestpath attribute entries using 3744 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
10 BGP extended community entries using 480 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 17264 total bytes of memory
BGP activity 101/84 prefixes, 2825/2791 paths, scan interval 60 secs
25 networks peaked at 14:54:41 Jan 26 2021 UTC (05:39:56.356 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001    5664    5668     5443    0    0 05:40:29       15
172.16.255.3    4        65001     378    5690     5443    0    0 05:35:23        5
172.16.255.4    4        65001     440    1633     5443    0    0 03:36:33        6
172.16.255.5    4        65001     594    5296     5443    0    0 04:34:27        8

Spine-01#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn
BGP table version is 5443, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.4:7
 *>i  [1][172.16.254.4:7][00010101010101010101][4294967295]/23
                       172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [1][172.16.254.4:101][00010101010101010101][0]/23
                       172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:9
```

```
 *>i  [1][172.16.254.5:9][00010101010101010101][4294967295]/23
                     172.16.254.5             0    100      0 ?
 * i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 *>i  [1][172.16.254.5:101][00010101010101010101][0]/23
                     172.16.254.5             0    100      0 ?
 * i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                     172.16.254.3             0    100      0 ?
 *>i                 172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                     172.16.254.3             0    100      0 ?
 *>i                 172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                     172.16.254.3             0    100      0 ?
 *>i                 172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                     172.16.254.4             0    100      0 ?
 *>i                 172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [2][172.16.254.5:101][0][48][44D3CA286C82][0][*]/20
                     172.16.254.5             0    100      0 ?
 *>i                 172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][44D3CA286CC2][0][*]/20
                     172.16.254.5             0    100      0 ?
 *>i                 172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                     172.16.254.5             0    100      0 ?
 *>i                 172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                     172.16.254.3             0    100      0 ?
 *>i                 172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
                     172.16.254.4             0    100      0 ?
 *>i                 172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [3][172.16.254.5:101][0][32][172.16.254.5]/17
                     172.16.254.5             0    100      0 ?
 *>i                 172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:257
 * i  [4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23
                     172.16.254.4             0    100      0 ?
 *>i                 172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:257
 * i  [4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23
                     172.16.254.5             0    100      0 ?
 *>i                 172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                     172.16.254.5             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i                 172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?

Spine-01#
```

The following example shows the output for the **show ip route** command on Spine Switch 1:

```
Spine-01# 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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 17 subnets, 2 masks
C        172.16.13.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.13.1/32 is directly connected, GigabitEthernet1/0/1
C        172.16.14.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.14.1/32 is directly connected, GigabitEthernet1/0/2
C        172.16.15.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.15.1/32 is directly connected, GigabitEthernet1/0/3
O        172.16.23.0/24
           [110/2] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.24.0/24
           [110/2] via 172.16.14.4, 03:37:00, GigabitEthernet1/0/2
O        172.16.25.0/24
           [110/2] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
O        172.16.254.3/32
           [110/2] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.254.4/32
           [110/2] via 172.16.14.4, 03:36:50, GigabitEthernet1/0/2
O        172.16.254.5/32
           [110/2] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
C        172.16.255.1/32 is directly connected, Loopback0
O        172.16.255.2/32
           [110/3] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
           [110/3] via 172.16.14.4, 03:37:00, GigabitEthernet1/0/2
           [110/3] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.255.3/32
           [110/2] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.255.4/32
           [110/2] via 172.16.14.4, 03:36:56, GigabitEthernet1/0/2
O        172.16.255.5/32
           [110/2] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
Spine-01#
```

Return to Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric, on page 583.

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 5499, main routing table version 5499
17 network entries using 5848 bytes of memory
34 path entries using 7072 bytes of memory
13/11 BGP path/bestpath attribute entries using 3744 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
```

```
10 BGP extended community entries using 480 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 17264 total bytes of memory
BGP activity 101/84 prefixes, 2823/2789 paths, scan interval 60 secs
25 networks peaked at 14:56:03 Jan 26 2021 UTC (05:40:54.652 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001    5669    5665     5499    0    0 05:41:28        15
172.16.255.3    4        65001     381    5691     5499    0    0 05:36:22         5
172.16.255.4    4        65001     440    1632     5499    0    0 03:37:31         6
172.16.255.5    4        65001     594    5291     5499    0    0 04:35:26         8

Spine-02#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn
BGP table version is 5499, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.4:7
 *>i  [1][172.16.254.4:7][00010101010101010101][4294967295]/23
                      172.16.254.4             0    100      0 ?
 * i                    172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [1][172.16.254.4:101][00010101010101010101][0]/23
                      172.16.254.4             0    100      0 ?
 * i                    172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:9
 *>i  [1][172.16.254.5:9][00010101010101010101][4294967295]/23
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 *>i  [1][172.16.254.5:101][00010101010101010101][0]/23
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i                    172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [2][172.16.254.5:101][0][48][44D3CA286C82][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                    172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                    172.16.254.5             0    100      0 ?
```

```
 * i  [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                     172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [3][172.16.254.5:101][0][32][172.16.254.5]/17
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:257
 * i  [4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:257
 * i  [4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                     172.16.254.5             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?

Spine-02#
```

The following example shows the output for the **show ip route** command on Spine Switch 2:

```
Spine-02# 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, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       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
       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR

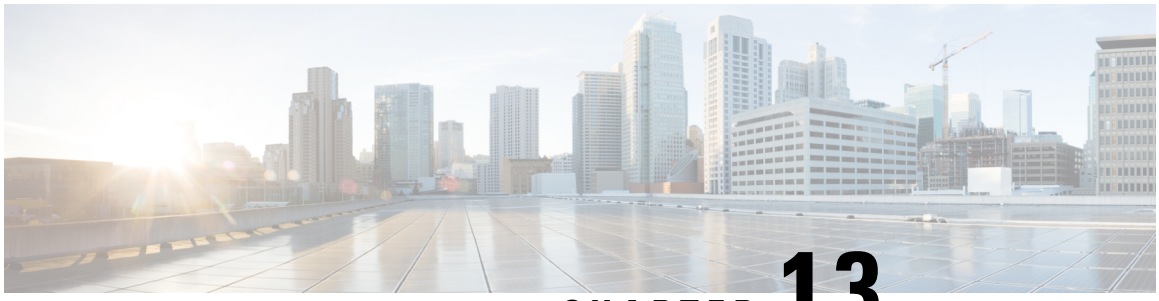
Gateway of last resort is not set

      172.16.0.0/16 is variably subnetted, 17 subnets, 2 masks
O        172.16.13.0/24
           [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
O        172.16.14.0/24
           [110/2] via 172.16.24.4, 03:37:38, GigabitEthernet1/0/2
O        172.16.15.0/24
           [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
C        172.16.23.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.23.2/32 is directly connected, GigabitEthernet1/0/1
C        172.16.24.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.24.2/32 is directly connected, GigabitEthernet1/0/2
C        172.16.25.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.25.2/32 is directly connected, GigabitEthernet1/0/3
O        172.16.254.3/32
```

```
          [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
O        172.16.254.4/32
          [110/2] via 172.16.24.4, 03:37:28, GigabitEthernet1/0/2
O        172.16.254.5/32
          [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
O        172.16.255.1/32
          [110/3] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
          [110/3] via 172.16.24.4, 03:37:38, GigabitEthernet1/0/2
          [110/3] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
C        172.16.255.2/32 is directly connected, Loopback0
O        172.16.255.3/32
          [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
O        172.16.255.4/32
          [110/2] via 172.16.24.4, 03:37:34, GigabitEthernet1/0/2
O        172.16.255.5/32
          [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3

Spine-02#
```

Return to Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric, on page 583.



CHAPTER 13

Configuring Private VLANs in a BGP EVPN VXLAN Fabric

- Restrictions for Private VLANs in a BGP EVPN VXLAN Fabric, on page 599
- Information About Private VLANs in a BGP EVPN VXLAN Fabric, on page 599
- How to Configure Private VLANs in a BGP EVPN VXLAN Fabric, on page 605
- Configuration Examples for Private VLANs in a BGP EVPN VXLAN Fabric, on page 608

Restrictions for Private VLANs in a BGP EVPN VXLAN Fabric

Configuration of Private VLANs in a BGP EVPN VXLAN fabric must be done in the following order:

1. Configure VLAN with primary and secondary associations.
2. Enable EVPN separately in each of the primary, community, and isolated VLANs.

   For more information, see Configuring an EVPN Instance on the VLAN on a VTEP, on page 20 .

If there's an EVPN configuration already associated with a VLAN, you can't directly configure the PVLAN associations for this VLAN. First disassociate the EVPN configuration from the VLAN. Next, configure the PVLAN association. Then reconfigure EVPN in each of the newly configured primary, community, and isolated VLANs.

Information About Private VLANs in a BGP EVPN VXLAN Fabric

A private VLAN (PVLAN) divides a regular VLAN into logical partitions, allowing limited broadcast boundaries among selected port-groups on a single Layer 2 Ethernet switch. The single Ethernet switch's PVLAN capabilities can be extended over the BGP EVPN VXLAN enabled network to build partitioned bridge-domain between port-groups across multiple Ethernet switches in the BGP EVPN VXLAN VTEP mode. The integration of PVLAN with a BGP EVPN VXLAN network enables the following benefits:

- Microsegmented Layer 2 network segregation across one or more BGP EVPN VXLAN switches.
- Partitioned and secured user-group Layer 2 network that limits the communication with dynamic or static port configuration assignments.
- IP subnet pool conservation across BGP EVPN VXLAN network while extending segregated Layer 2 network across the fabric.

- Conservation of Layer 2 overlay tunnels and peer networks with a single virtual network identifier (VNI) mapped to Primary VLAN.

Primary and Secondary VLANs

Each subdomain in a PVLAN is represented by a pair of VLANs: a primary VLAN and a secondary VLAN. A PVLAN can have multiple VLAN pairs, one pair for each subdomain. All VLAN pairs in a PVLAN share the same primary VLAN. The secondary VLAN ID differentiates one subdomain from another. A secondary VLAN can either be an isolated VLAN or a community VLAN. Primary and secondary VLANs have the following characteristics:

- Primary VLAN: A PVLAN has only one primary VLAN. Every port in a PVLAN is a member of the primary VLAN. The primary VLAN carries unidirectional traffic downstream from the promiscuous ports to the host (isolated and community) ports and to other promiscuous ports.
- Isolated VLAN: A PVLAN has only one isolated VLAN. An isolated VLAN is a secondary VLAN that carries unidirectional traffic upstream from the hosts towards the promiscuous ports and the gateway.
- Community VLAN: A community VLAN is a secondary VLAN that carries upstream traffic from the community ports to the promiscuous port gateways and to other host ports in the same community. You can configure multiple community VLANs in a PVLAN.

Private VLAN Ports

PVLAN ports are access ports that are one of these types:

- Promiscuous: A promiscuous port belongs to the primary VLAN. It can communicate with all interfaces, including the community and isolated host ports that belong to the secondary VLANs associated with the primary VLAN.
- Isolated: An isolated port is a host port that belongs to an isolated secondary VLAN. It has complete Layer 2 separation from other ports within the same PVLAN, except for the promiscuous ports. PVLANs block all traffic to isolated ports except traffic from promiscuous ports. Likewise, PVLANs forward the traffic from an isolated port only to promiscuous ports.
- Community: A community port is a host port that belongs to a community secondary VLAN. Community ports communicate with other ports in the same community VLAN and with promiscuous ports. Community ports are isolated at Layer 2 from all other interfaces in external communities and also from isolated ports within their private VLAN.

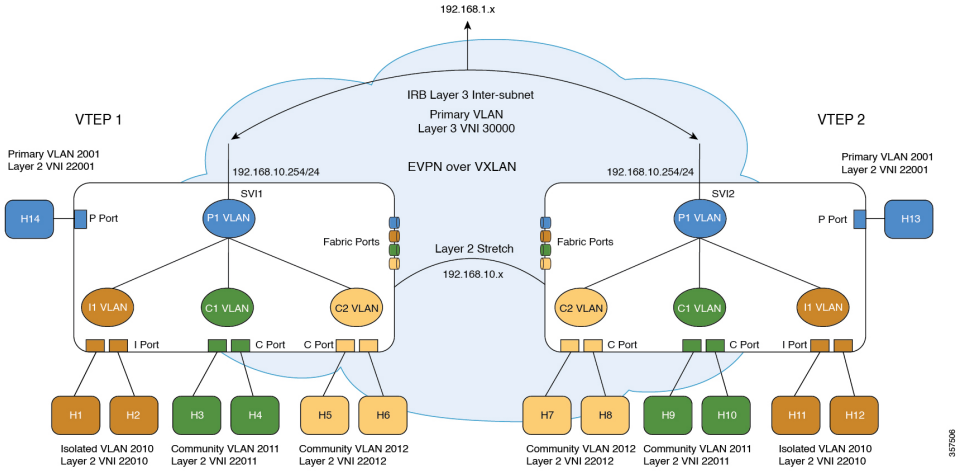
Extension of Private VLANs in a BGP EVPN VXLAN Fabric

Private VLANs (PVLANs) partition a regular VLAN domain into subdomains and provide Layer 2 isolation between ports within the same PVLAN. Like a regular VLAN, a private VLAN can span multiple Layer 2 switches. In a private VLAN that spans across multiple devices, traffic from an isolated port in Switch A does not reach an isolated port on Switch B. This is achieved by the trunk port carrying the primary VLAN and secondary VLANs to neighboring switches with dot1q tag in a traditional Layer 2 network. With BGP EVPN VXLAN enabled in the PVLANs on the VTEPs, the L2VNI segment preserves the PVLAN semantics and provides the Layer 2 isolation for the stretched PVLAN segment across the VTEPs in the overlay fabric. PVLAN extension with BGP EVPN VXLAN allows you to:

- Seamlessly migrate and join (or stretch) the PVLAN domain like any regular VLAN.
- Access to centralized common services such as printer or DHCP through the promiscuous port on any VTEP in the EVPN overlay.
- Maintain community and isolated VLAN semantics in the overlay fabric across all the VTEPs. The EVPN fabric provides a logical single switch view for the respective Layer 2 domain.

The following image shows PVLAN extension in a BGP EVPN VXLAN fabric with two VTEPs:

Figure 51: PVLAN Extension in a BGP EVPN VXLAN Fabric



Traffic Forwarding for Private VLANs in a BGP EVPN VXLAN Fabric

You can forward known unicast and broadcast, unknown unicast, and multicast (BUM) traffic between PVLANs in a BGP EVPN VXLAN fabric. On the Source VTEP, the forwarding process on the access PVLAN ports (promiscuous, isolated, community) adheres to the baseline PVLAN forwarding. With BGP EVPN VXLAN enabled in the PVLAN domain, the remote host routes are learned and programmed in the hardware of the respective PVLANs. The following sections illustrate the forwarding scenarios for unicast and BUM traffic between local and remote hosts for each of the secondary VLANs.

Known Unicast Traffic Forwarding

The sending VTEP bridges a known unicast packet with the corresponding secondary VLAN's virtual network identifier (VNI) ID. The packet arrives on the receiving VTEP. After decapsulation, receiving VTEP processes the packet in the same way as a packet from a local PVLAN host port. The packet gets mapped to the respective community, isolated, or primary VLAN.

The following images illustrate the known unicast traffic forwarding scenarios for PVLANs in a BGP EVPN VXLAN fabric:

Figure 52: Unicast Traffic from Promiscuous Port : H14 to H8

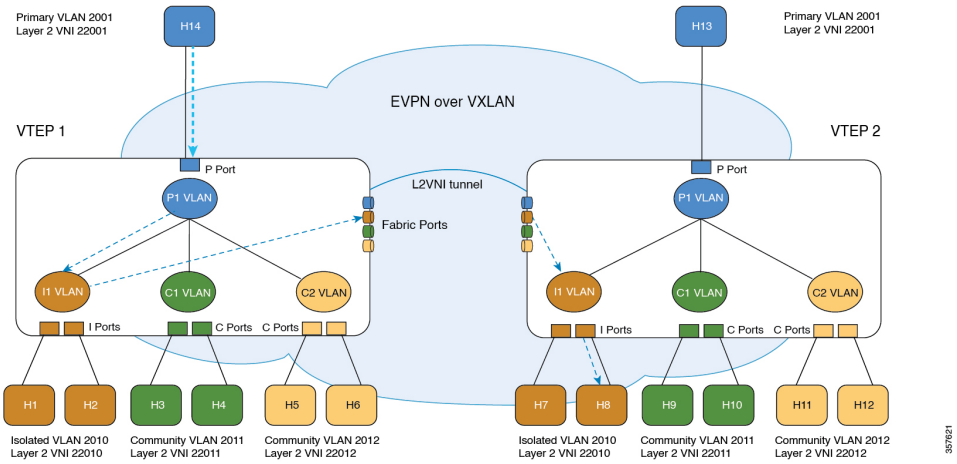


Figure 53: Unicast Traffic from Isolated Port: H1 to H13

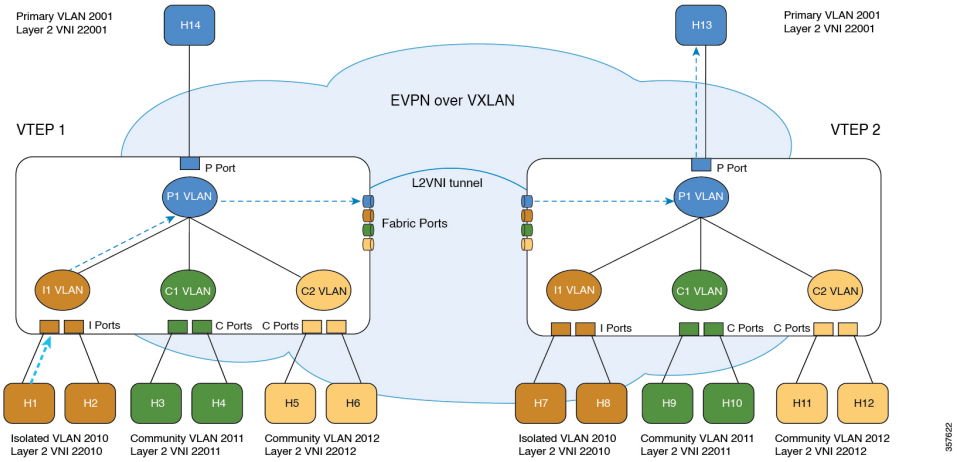
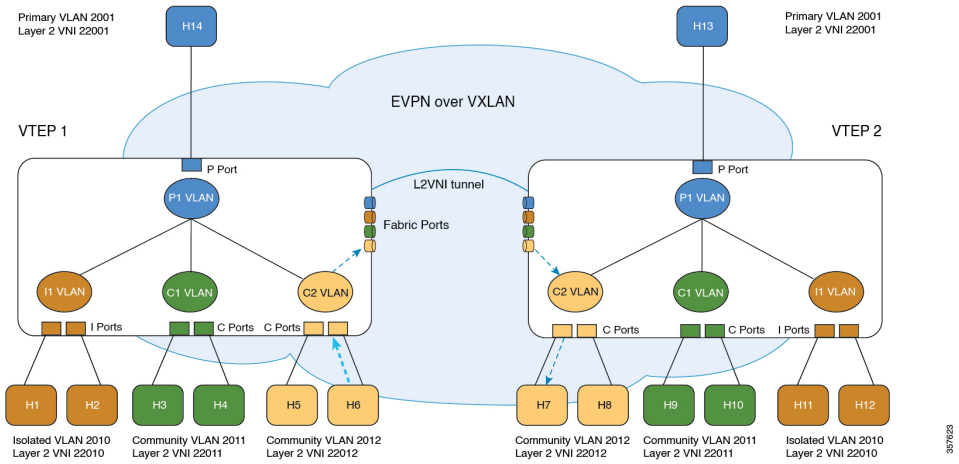


Figure 54: Unicast Traffic from Community Port: H6 to H7



Broadcast, Unknown Unicast, and Multicast Traffic Forwarding

In a regular VLAN, broadcasts are forwarded to all ports in that VLAN. Private VLAN broadcast forwarding depends on the port sending the broadcast:

- An isolated port sends a broadcast only to the promiscuous ports or trunk ports.
- A community port sends a broadcast to all promiscuous ports, trunk ports, and ports in the same community VLAN.
- A promiscuous port sends a broadcast to all ports in the private VLAN (other promiscuous ports, trunk ports, isolated ports, and community ports).

In addition to the above, a copy of the flood packet is sent to the remote VTEPs with the respective L2VNI. (See BUM traffic handling in the L2VNI). On the remote VTEP, the flood copy is again replicated towards the access as per the PVLAN broadcast rules mentioned above. Flood packets received from the fabric are not sent back to fabric with split-horizon check.

During forwarding, if a packet's MAC address isn't available in the lookup, the VTEP replicates the packet with the VNI ID of the forwarding (or incoming) VLAN. The VTEP forwards the BUM packets with the VNI ID of the corresponding VLAN. The receiving VTEP decapsulates the BUM packet and maps the VNI ID to the corresponding secondary VLAN. This mapping ensures that the flood rules remain local. The VTEP then processes the packet in the same way as a packet from a local host port.

For isolated VLANs, after the destination MAC address lookup results in an unknown unicast from the source port, it's not locally known whether the destination MAC address belongs to the remote isolated VLAN host or the remote primary VLAN host. Hence, the BUM packet copy is allowed to go the egress VTEPs with the isolated VLAN VNI ID. On egress VTEPs, this BUM copy gets flooded on local isolated ports and local promiscuous ports. As a result, BUM traffic from remote isolated ports to local isolated ports is unavoidable.

Note Forwarding of unknown unicast traffic from an isolated port to a remote promiscuous port isn't supported.

The following images illustrate the BUM traffic forwarding scenarios for PVLANs in a BGP EVPN VXLAN fabric:

Figure 55: BUM Traffic from Promiscuous Port (H14)

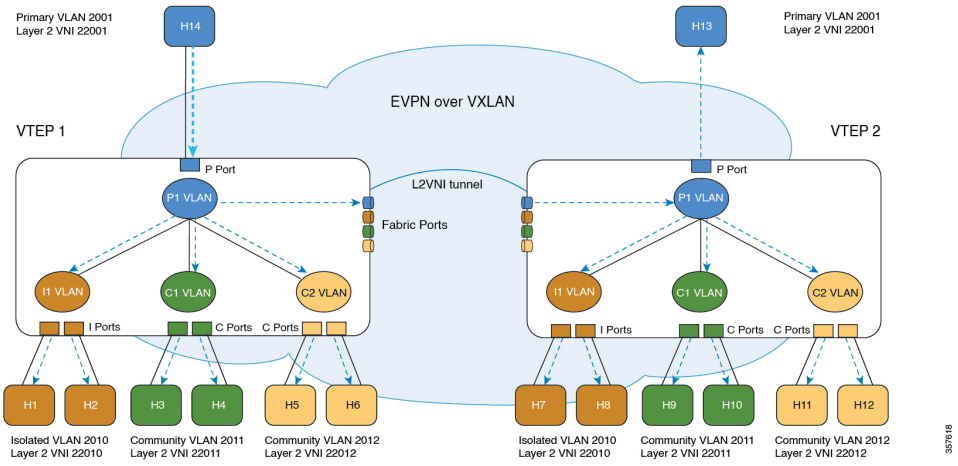


Figure 56: BUM Traffic from Isolated Port (H1)

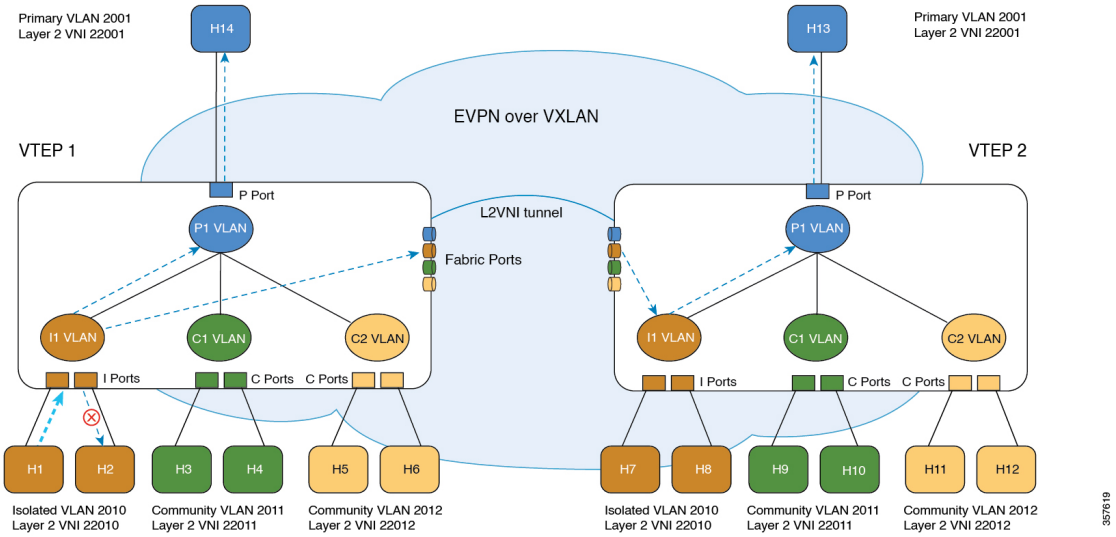
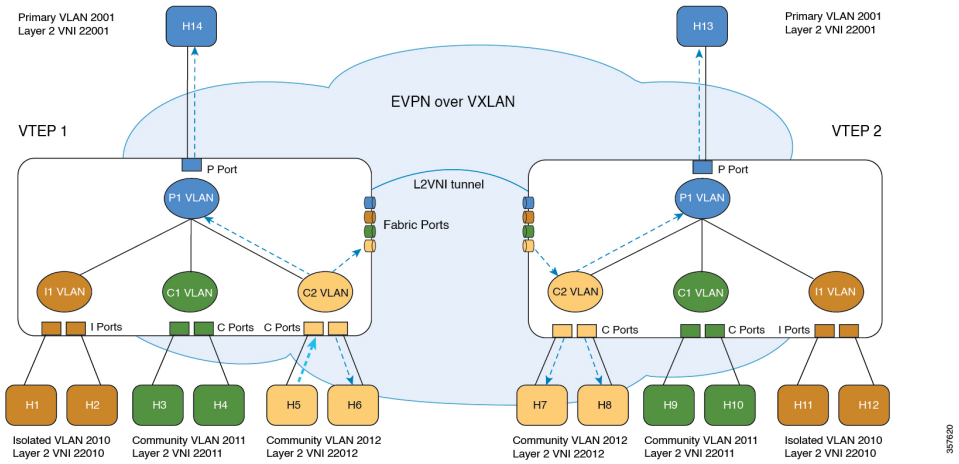


Figure 57: BUM Traffic from Community Port (H5)



Routed Traffic Forwarding

Routed traffic between the hosts in a microsegmented VLAN is through the associated Primary VLAN SVI on the Local VTEP (For more information, see "Configuring Private VLANs" module in the *VLAN Configuration Guide* for the applicable release). When the source and destination hosts are across the EVPN VXLAN fabric, the routed traffic between the microsegmented VLAN hosts follows the Symmetric Integrated Routing and Bridging (IRB) method to cross the fabric (For more information, see Configuring EVPN VXLAN Integrated Routing and Bridging, on page 103). On the destination VTEP, traffic is routed from the core VLAN SVI to the associated Primary VLAN SVI interface and then bridged in the microsegmented local destination Secondary VLAN.

How to Configure Private VLANs in a BGP EVPN VXLAN Fabric

When you configure PVLANs in a BGP EVPN VXLAN fabric, the existing PVLAN configuration is preserved and the Layer 2 VNI configuration is added to the PVLAN. By adding the Layer 2 VNI configuration, you expand the PVLAN and stretch it over the fabric across the VTEPs in the fabric.

In a BGP EVPN VXLAN fabric, the EVPN control plane distributes the MAC and MAC-IP routes. In addition, PVLANs handle BUM and unicast traffic forwarding differently compared to regular VLANs. Due to these two reasons, you can create and delete PVLANs strictly in the following ways:

- To create a PVLAN, first configure the VLAN with primary and secondary associations. Next, enable EVPN separately in each of the primary, community, and isolated VLANs.

Note If there's an EVPN configuration already associated with a VLAN, you can't directly configure the PVLAN associations for this VLAN. First, use the **member vni** command in VLAN configuration mode to disassociate the EVPN configuration from the VLAN. Next, configure the PVLAN association. Now reconfigure EVPN in each of the newly configured primary, community, and isolated VLANs.

- To delete a PVLAN, ensure that you unconfigure EVPN in the respective VLAN before you modify the PVLAN configuration.

Configuring the Primary and Secondary VLANs for a Private VLAN

To configure the primary and secondary VLANs for a private VLAN, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan** *vlan-id* **Example:** `Device(config)# vlan 101`	Enters VLAN configuration mode for the specified VLAN ID.

Step 4	**private-vlan** {**association** [**add** | **remove**] *secondary-vlan-list* | **community** | **isolated** | **primary**}
	Example:
	`Device(config-vlan)# private-vlan primary`
	`Device(config-vlan)# private-vlan association 102`
		Configures the VLAN as a PVLAN and configures the association between primary and secondary VLANs.
		Use the **primary** keyword to configure the VLAN as a PVLAN.
		Use the **community** keyword to designate the VLAN as a community VLAN.
		Use the **isolated** keyword to designate the VLAN as an isolated VLAN.
		Use the **association** [**add** | **remove**] keyword to add or remove the association between a primary and secondary VLAN.

Step 5	**exit** **Example:** `Device(config-vlan)# exit`	Exits VLAN configuration mode and returns to global configuration mode.
Step 6	Repeat steps 3 to 5 for each primary and secondary VLAN, as needed.	--
Step 7	**end** **Example:** `Device(config)# end`	Exits global configuration mode and enters privileged EXEC mode.

Configuring the Port of a Private VLAN

To configure a port of a PVLAN, perform the following steps:

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface** *interface-id* **Example:** `Device(config)# interface GigabitEthernet1/0/1`	Enters interface configuration mode for the specified interface ID.
Step 4	**switchport mode private-vlan** {**host** | **promiscuous**} **Example:** `Device(config-if)# switchport mode private-vlan host`	Configures the interface as either a host PVLAN port or a promiscuous PVLAN port.
Step 5	**switchport private-vlan** {**host-association** | **mapping** *primary-vlan-id secondary-vlan-id-list* **Example:** `Device(config-if)# switchport private-vlan host-association 101 104`	Associates a PVLAN host port or maps a PVLAN promiscuous port to a primary VLAN. **Note** If you configure a port as a PVLAN host port and you do not configure a valid PVLAN association with the **switchport private-vlan host-association** command, the interface becomes inactive. **Note** If you configure a port as a PVLAN promiscuous port and you do not configure a valid PVLAN mapping with the **switchport private-vlan mapping** command, the interface becomes inactive.
Step 6	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and enters privileged EXEC mode.

Enabling EVPN in a Private VLAN

To enable EVPN in a PVLAN, perform the following steps:

Note Enable EVPN separately in each of the primary, community, and isolated VLANs.

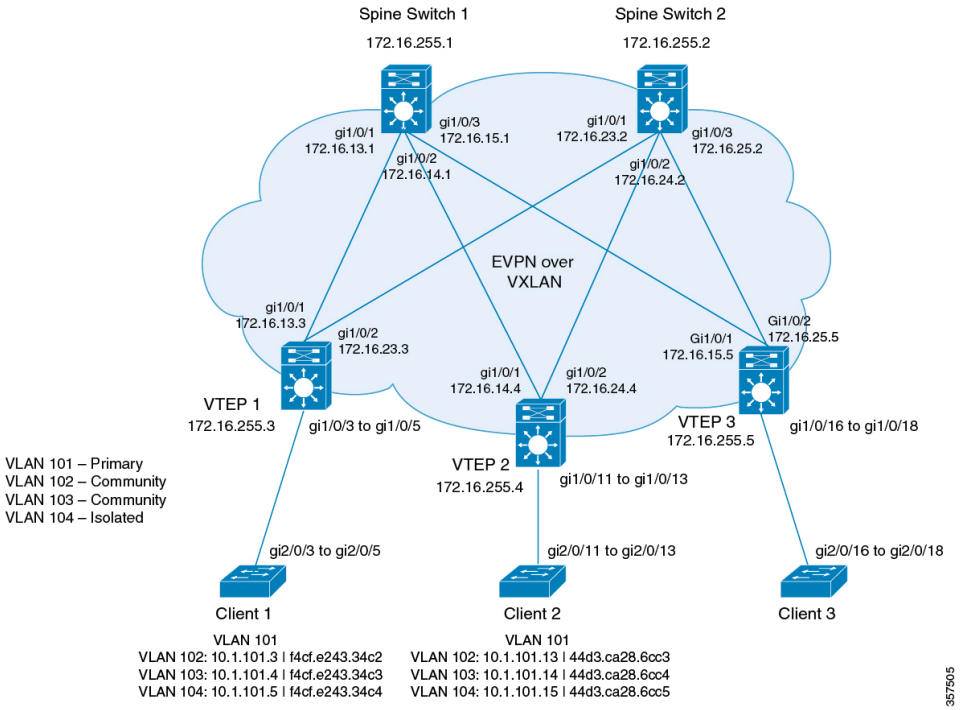
Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan configuration** *vlan-id* **Example:** `Device(config)# vlan configuration 101`	Enters VLAN configuration mode for the specified PVLAN interface.
Step 4	**member evpn-instance** *evpn-instance-id* **vni** *layer2-vni-id* **Example:** `Device(config-vlan)# member evpn-instance 1 vni 6000`	Adds EVPN instance as a member of the PVLAN configuration. The VNI here is used as a Layer 2 VNI.
Step 5	**end** **Example:** `Device(config-vlan)# end`	Exits VLAN configuration mode and enters privileged EXEC mode.

Configuration Examples for Private VLANs in a BGP EVPN VXLAN Fabric

This section provides a configuration example for PVLANs in a BGP EVPN VXLAN fabric using the following topology:

Figure 58: Private VLANs in a BGP EVPN VXLAN Fabric



The topology shows an EVPN VXLAN network with two spine switches (Spine Switch 1 and Spine Switch 2) and three VTEPs (VTEP 1, VTEP 2, and VTEP 3). The network has an extended PVLAN with VLAN 101 as the primary VLAN. VLAN 102, VLAN 103, and VLAN 104 are the secondary VLANs. The following tables provide the sample configurations for the devices in this topology:

Table 55: Configuring VTEP 1, VTEP 2, and VTEP 3 for PVLAN Extension in a BGP EVPN VXLAN Fabric

VTEP 1	VTEP 2	VTEP 3
Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ! vtp mode transparent ! l2vpn evpn replication-type static default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! l2vpn evpn instance 103 vlan-based encapsulation vxlan ! l2vpn evpn instance 104 vlan-based encapsulation vxlan ! l2vpn evpn instance 201 vlan-based encapsulation vxlan ! l2vpn evpn instance 202 vlan-based encapsulation vxlan	Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ! vtp mode transparent ! l2vpn evpn replication-type static default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! l2vpn evpn instance 103 vlan-based encapsulation vxlan ! l2vpn evpn instance 104 vlan-based encapsulation vxlan ! l2vpn evpn instance 201 vlan-based encapsulation vxlan ! l2vpn evpn instance 202 vlan-based encapsulation vxlan	Leaf-03# **show running-config** hostname Leaf-03 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! ip multicast-routing ! vtp mode transparent ! l2vpn evpn replication-type static default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan ! l2vpn evpn instance 103 vlan-based encapsulation vxlan ! l2vpn evpn instance 104 vlan-based encapsulation vxlan ! l2vpn evpn instance 201 vlan-based encapsulation vxlan ! l2vpn evpn instance 202 vlan-based encapsulation vxlan

VTEP 1	VTEP 2	VTEP 3
!	!	!
l2vpn evpn instance 203	l2vpn evpn instance 203	l2vpn evpn instance 203
vlan-based	vlan-based	vlan-based
encapsulation vxlan	encapsulation vxlan	encapsulation vxlan
!	!	!
l2vpn evpn instance 204	l2vpn evpn instance 204	l2vpn evpn instance 204
vlan-based	vlan-based	vlan-based
encapsulation vxlan	encapsulation vxlan	encapsulation vxlan
!	!	!
system mtu 9198	system mtu 9198	system mtu 9198
!	!	!
vlan configuration 101	vlan configuration 101	vlan configuration 101
member evpn-instance 101 vni 10101	member evpn-instance 101 vni 10101	member evpn-instance 101 vni 10101
vlan configuration 102	vlan configuration 102	vlan configuration 102
member evpn-instance 102 vni 10102	member evpn-instance 102 vni 10102	member evpn-instance 102 vni 10102
vlan configuration 103	vlan configuration 103	vlan configuration 103
member evpn-instance 103 vni 10103	member evpn-instance 103 vni 10103	member evpn-instance 103 vni 10103
vlan configuration 104	vlan configuration 104	vlan configuration 104
member evpn-instance 104 vni 10104	member evpn-instance 104 vni 10104	member evpn-instance 104 vni 10104
vlan configuration 201	vlan configuration 201	vlan configuration 201
member evpn-instance 201 vni 10201	member evpn-instance 201 vni 10201	member evpn-instance 201 vni 10201
vlan configuration 202	vlan configuration 202	vlan configuration 202
member evpn-instance 202 vni 10202	member evpn-instance 202 vni 10202	member evpn-instance 202 vni 10202
vlan configuration 203	vlan configuration 203	vlan configuration 203
member evpn-instance 203 vni 10203	member evpn-instance 203 vni 10203	member evpn-instance 203 vni 10203
vlan configuration 204	vlan configuration 204	vlan configuration 204
member evpn-instance 204 vni 10204	member evpn-instance 204 vni 10204	member evpn-instance 204 vni 10204
vlan configuration 901	vlan configuration 901	vlan configuration 901
member vni 50901	member vni 50901	member vni 50901
!	!	!
vlan 101	vlan 101	vlan 101
private-vlan primary	private-vlan primary	private-vlan primary
private-vlan association 102-104	private-vlan association 102-104	private-vlan association 102-104
!	!	!
vlan 102	vlan 102	vlan 102
private-vlan community	private-vlan community	private-vlan community
!	!	!
vlan 103	vlan 103	vlan 103
private-vlan community	private-vlan community	private-vlan community
!	!	!
vlan 104	vlan 104	vlan 104
private-vlan isolated	private-vlan isolated	private-vlan isolated
!	!	!
vlan 201	vlan 201	vlan 201
private-vlan primary	private-vlan primary	private-vlan primary
private-vlan association 202-204	private-vlan association 202-204	private-vlan association 202-204
!	!	!
vlan 202	vlan 202	vlan 202
private-vlan community	private-vlan community	private-vlan community
!	!	!

VTEP 1	VTEP 2	VTEP 3
vlan 203 private-vlan community ! vlan 204 private-vlan isolated ! vlan 901 ! interface Loopback0 ip address 172.16.255.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.13.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.23.3 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/3 switchport access vlan 102 switchport private-vlan host-association 101 102 switchport mode private-vlan host spanning-tree portfast ! interface GigabitEthernet1/0/4 switchport access vlan 103 switchport private-vlan host-association 101 103 switchport mode private-vlan host spanning-tree portfast !	vlan 203 private-vlan community ! vlan 204 private-vlan isolated ! vlan 901 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.4 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/11 switchport access vlan 102 switchport private-vlan host-association 101 102 switchport mode private-vlan host spanning-tree portfast ! interface GigabitEthernet1/0/12 switchport access vlan 103 switchport private-vlan host-association 101 103 switchport mode private-vlan host spanning-tree portfast !	vlan 203 private-vlan community ! vlan 204 private-vlan isolated ! vlan 901 ! interface Loopback0 ip address 172.16.255.5 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.5 255.255.255.255 ip pim sparse-mode ip ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.15.5 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.25.5 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/16 switchport access vlan 202 switchport private-vlan host-association 201 202 switchport mode private-vlan host spanning-tree portfast !

VTEP 1	VTEP 2	VTEP 3
interface GigabitEthernet1/0/5 switchport access vlan 104 switchport private-vlan host-association 101 104 switchport mode private-vlan host spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 private-vlan mapping 102-104 ! interface Vlan201 vrf forwarding green ip address 10.1.201.1 255.255.255.0 private-vlan mapping 202-204 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.1.1.1 member vni 10102 mcast-group 225.1.1.1 member vni 10103 mcast-group 225.1.1.1 member vni 10104 mcast-group 225.1.1.1 member vni 10201 mcast-group 225.1.1.1 member vni 10202 mcast-group 225.1.1.1 member vni 10203 mcast-group 225.1.1.1 member vni 10204 mcast-group 225.1.1.1 member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.3 !	interface GigabitEthernet1/0/13 switchport access vlan 104 switchport private-vlan host-association 101 104 switchport mode private-vlan host spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 private-vlan mapping 102-104 ! interface Vlan201 vrf forwarding green ip address 10.1.201.1 255.255.255.0 private-vlan mapping 202-204 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.1.1.1 member vni 10102 mcast-group 225.1.1.1 member vni 10103 mcast-group 225.1.1.1 member vni 10104 mcast-group 225.1.1.1 member vni 10201 mcast-group 225.1.1.1 member vni 10202 mcast-group 225.1.1.1 member vni 10203 mcast-group 225.1.1.1 member vni 10204 mcast-group 225.1.1.1 member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.4 !	interface GigabitEthernet1/0/17 switchport access vlan 203 switchport private-vlan host-association 201 203 switchport mode private-vlan host spanning-tree portfast ! interface GigabitEthernet1/0/18 switchport access vlan 204 switchport private-vlan host-association 201 204 switchport mode private-vlan host spanning-tree portfast ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 private-vlan mapping 102-104 ! interface Vlan201 vrf forwarding green ip address 10.1.201.1 255.255.255.0 private-vlan mapping 202-204 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 10101 mcast-group 225.1.1.1 member vni 10102 mcast-group 225.1.1.1 member vni 10103 mcast-group 225.1.1.1 member vni 10104 mcast-group 225.1.1.1 member vni 10201 mcast-group 225.1.1.1 member vni 10202 mcast-group 225.1.1.1 member vni 10203 mcast-group 225.1.1.1 member vni 10204 mcast-group 225.1.1.1 member vni 50901 vrf green !

VTEP 1	VTEP 2	VTEP 3
router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-01#	router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-02#	router ospf 1 router-id 172.16.255.5 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-03#

Table 56: Configuring Spine Switch 1 and Spine Switch 2 for PVLAN Extension in a BGP EVPN VXLAN Fabric

Spine Switch 1

```
Spine-01# show running-config
hostname Spine-01
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.1 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.13.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.14.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/3
no switchport
ip address 172.16.15.1 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.1
!
router bgp 65001
bgp router-id 172.16.255.1
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.2 remote-as 65001
neighbor 172.16.255.2 update-source Loopback0
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.5 remote-as 65001
neighbor 172.16.255.5 update-source Loopback0
!
```

Spine Switch 2

```
Spine-02# show running-config
hostname Spine-02
!
ip routing
!
ip multicast-routing
!
system mtu 9198
!
interface Loopback0
ip address 172.16.255.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback1
ip address 172.16.254.2 255.255.255.255
ip ospf 1 area 0
!
interface Loopback2
ip address 172.16.255.255 255.255.255.255
ip pim sparse-mode
ip ospf 1 area 0
!
interface GigabitEthernet1/0/1
no switchport
ip address 172.16.23.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/2
no switchport
ip address 172.16.24.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
interface GigabitEthernet1/0/3
no switchport
ip address 172.16.25.2 255.255.255.0
ip pim sparse-mode
ip ospf network point-to-point
ip ospf 1 area 0
!
router ospf 1
router-id 172.16.255.2
!
router bgp 65001
bgp router-id 172.16.255.2
bgp log-neighbor-changes
no bgp default ipv4-unicast
neighbor 172.16.255.1 remote-as 65001
neighbor 172.16.255.1 update-source Loopback0
neighbor 172.16.255.3 remote-as 65001
neighbor 172.16.255.3 update-source Loopback0
neighbor 172.16.255.4 remote-as 65001
neighbor 172.16.255.4 update-source Loopback0
neighbor 172.16.255.5 remote-as 65001
neighbor 172.16.255.5 update-source Loopback0
!
```

Spine Switch 1	Spine Switch 2
address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.2 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.5 activate neighbor 172.16.255.5 send-community both neighbor 172.16.255.5 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.2 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end Spine-01#	address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate neighbor 172.16.255.1 send-community both neighbor 172.16.255.1 route-reflector-client neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-reflector-client neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-reflector-client neighbor 172.16.255.5 activate neighbor 172.16.255.5 send-community both neighbor 172.16.255.5 route-reflector-client exit-address-family ! ip pim rp-address 172.16.255.255 ip msdp peer 172.16.254.1 connect-source Loopback1 remote-as 65001 ip msdp cache-sa-state ! end Spine-02#

Outputs to Verify the Configuration on VTEP 1

The following example shows the output for the **show vlan private-vlan** command on VTEP 1:

```
Leaf-01# show vlan private-vlan

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
101     102       community         Gi1/0/3
101     103       community         Gi1/0/4
101     104       isolated          Gi1/0/5
201     202       community
201     203       community
201     204       isolated

Leaf-01#
```

The following example shows the output for the **show ip arp vrf green** command on VTEP 1:

```
Leaf-01# show ip arp vrf green
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.101.1              -   10b3.d56a.8fc1  ARPA   Vlan101
Internet  10.1.101.3             95   f4cf.e243.34c2  ARPA   Vlan101 pv 102
Internet  10.1.101.4             95   f4cf.e243.34c3  ARPA   Vlan101 pv 103
Internet  10.1.101.5             95   f4cf.e243.34c4  ARPA   Vlan101 pv 104
Internet  10.1.201.1              -   10b3.d56a.8fcc  ARPA   Vlan201
Internet  172.16.254.3            -   10b3.d56a.8fc8  ARPA   Vlan901

Leaf-01#
```

The following example shows the output for the **show mac address-table vlan** *vlan-id* command on VTEP 1:

```
Leaf-01# show mac address-table vlan 101
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
 101    10b3.d56a.8fc1    STATIC      Vl101
 101    7c21.0dbd.9541    STATIC      Vl101
 101    f4cf.e243.34c2    DYNAMIC pv  Gi1/0/3
 101    f4cf.e243.34c3    DYNAMIC pv  Gi1/0/4
 101    f4cf.e243.34c4    DYNAMIC pv  Gi1/0/5
Total Mac Addresses for this criterion: 5

Leaf-01#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 1:

```
Leaf-01# show l2vpn evpn peers vxlan

Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.4            8          10101    01:33:29
nve1      10102    172.16.254.4            1          10102    01:33:29
nve1      10103    172.16.254.4            1          10103    01:33:29
nve1      10104    172.16.254.4            1          10104    00:01:37

Leaf-01#
```

The following example shows the output for the **show nve peer** command on VTEP 1:

```
Leaf-01# show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.5     7c21.0dbd.2748 50901      UP  A/M/4 01:33:30
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/4 01:33:29
nve1       10101    L2CP 172.16.254.4     8              10101      UP   N/A  01:33:29
nve1       10102    L2CP 172.16.254.4     1              10102      UP   N/A  01:33:29
nve1       10103    L2CP 172.16.254.4     1              10103      UP   N/A  01:33:29
nve1       10104    L2CP 172.16.254.4     1              10104      UP   N/A  00:01:37

Leaf-01#
```

The following example shows the output for the **show l2vpn evpn mac local** command on VTEP 1:

```
Leaf-01# show l2vpn evpn mac local
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
f4cf.e243.34c2 101   101   0000.0000.0000.0000.0000 0          Gi1/0/3
f4cf.e243.34c3 101   101   0000.0000.0000.0000.0000 0          Gi1/0/4
f4cf.e243.34c4 101   101   0000.0000.0000.0000.0000 0          Gi1/0/5
f4cf.e243.34c2 102   102   0000.0000.0000.0000.0000 0          Gi1/0/3
f4cf.e243.34c3 103   103   0000.0000.0000.0000.0000 0          Gi1/0/4
f4cf.e243.34c4 104   104   0000.0000.0000.0000.0000 0          Gi1/0/5

Leaf-01#
```

The following example shows the output for the **show l2vpn evpn mac remote** command on VTEP 1:

```
Leaf-01# show l2vpn evpn mac remote
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc3 101   101   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc4 101   101   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc5 101   101   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc3 102   102   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc4 103   103   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc5 104   104   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc6 201   201   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc7 201   201   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc8 201   201   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc6 202   202   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc7 203   203   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc8 204   204   0000.0000.0000.0000.0000 0          172.16.254.5

Leaf-01#
```

The following example shows the output for the **show l2route evpn mac ip** command on VTEP 1:

```
Leaf-01# show l2route evpn mac ip
  EVI       ETag  Prod    Mac Address          Host IP                 Next Hop(s)
----- ---------- ----- -------------- --------------- --------------------------
  101          0 L2VPN 10b3.d56a.8fc1      10.1.101.1                     Vl101:0
  101          0   BGP 44d3.ca28.6cc3     10.1.101.13       V:10101 172.16.254.4
  101          0   BGP 44d3.ca28.6cc4     10.1.101.14       V:10101 172.16.254.4
  101          0   BGP 44d3.ca28.6cc5     10.1.101.15       V:10101 172.16.254.4
  101          0   BGP 7c21.0dbd.9541      10.1.101.1       V:10101 172.16.254.4
  101          0 L2VPN f4cf.e243.34c2      10.1.101.3                Gi1/0/3:101
  101          0 L2VPN f4cf.e243.34c3      10.1.101.4                Gi1/0/4:101
  101          0 L2VPN f4cf.e243.34c4      10.1.101.5                Gi1/0/5
  201          0   BGP 44d3.ca28.6cc6      10.1.102.3       V:10201 172.16.254.5
  201          0   BGP 44d3.ca28.6cc7      10.1.102.4       V:10201 172.16.254.5
  201          0   BGP 44d3.ca28.6cc8      10.1.102.5       V:10201 172.16.254.5
  201          0   BGP 7c21.0dbd.274c      10.1.201.1       V:10201 172.16.254.5

Leaf-01#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 1:

```
Leaf-01# show bgp l2vpn evpn
BGP table version is 70, local router ID is 172.16.255.3
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.255.3:101
 *>   [2][172.16.255.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                       ::                                    32768 ?
 *>i  [2][172.16.255.3:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                       172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.3:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                       172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.3:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                       172.16.254.4             0    100      0 ?
```

```
 *>i  [2][172.16.255.3:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                     172.16.254.4             0    100      0 ?
 *>   [2][172.16.255.3:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                     ::                                    32768 ?
 *>   [2][172.16.255.3:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
                     ::                                    32768 ?
 *>   [2][172.16.255.3:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                     ::                                    32768 ?
Route Distinguisher: 172.16.255.3:102
 *>i  [2][172.16.255.3:102][0][48][44D3CA286CC3][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>   [2][172.16.255.3:102][0][48][F4CFE24334C2][0][*]/20
                     ::                                    32768 ?
Route Distinguisher: 172.16.255.3:103
 *>i  [2][172.16.255.3:103][0][48][44D3CA286CC4][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>   [2][172.16.255.3:103][0][48][F4CFE24334C3][0][*]/20
                     ::                                    32768 ?
Route Distinguisher: 172.16.255.3:104
 *>i  [2][172.16.255.3:104][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>   [2][172.16.255.3:104][0][48][F4CFE24334C4][0][*]/20
                     ::                                    32768 ?
Route Distinguisher: 172.16.255.3:201
 *>i  [2][172.16.255.3:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                     172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.3:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                     172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.3:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                     172.16.254.5             0    100      0 ?
     Network          Next Hop            Metric LocPrf Weight Path
 *>i  [2][172.16.255.3:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                     172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.3:202
 *>i  [2][172.16.255.3:202][0][48][44D3CA286CC6][0][*]/20
                     172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.3:203
 *>i  [2][172.16.255.3:203][0][48][44D3CA286CC7][0][*]/20
                     172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.3:204
 *>i  [2][172.16.255.3:204][0][48][44D3CA286CC8][0][*]/20
                     172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:101
 *>i  [2][172.16.255.4:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:102
 *>i  [2][172.16.255.4:102][0][48][44D3CA286CC3][0][*]/20
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:103
 *>i  [2][172.16.255.4:103][0][48][44D3CA286CC4][0][*]/20
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:104
```

```
 * i  [2][172.16.255.4:104][0][48][44D3CA286CC5][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:201
 *>i  [2][172.16.255.5:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.5:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.5:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.5:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                      172.16.254.5             0    100      0 ?
     Network          Next Hop            Metric LocPrf Weight Path
 * i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:202
 *>i  [2][172.16.255.5:202][0][48][44D3CA286CC6][0][*]/20
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:203
 *>i  [2][172.16.255.5:203][0][48][44D3CA286CC7][0][*]/20
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:204
 *>i  [2][172.16.255.5:204][0][48][44D3CA286CC8][0][*]/20
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1 (default for vrf green)
 *>   [5][1:1][0][24][10.1.101.0]/17
                      0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][24][10.1.201.0]/17
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?

Leaf-01#
```

Outputs to Verify the Configuration on VTEP 2

The following example shows the output for the **show vlan private-vlan** command on VTEP 2:

```
Leaf-02# show vlan private-vlan
Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
101     102       community         Gi1/0/11
101     103       community         Gi1/0/12
101     104       isolated          Gi1/0/13
201     202       community
201     203       community
201     204       isolated

Leaf-02#
```

The following example shows the output for the **show ip arp vrf green** command on VTEP 2:

```
Leaf-02# show ip arp vrf green
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.101.1              -   7c21.0dbd.9541  ARPA   Vlan101
Internet  10.1.101.13            95   44d3.ca28.6cc3  ARPA   Vlan101 pv 102
Internet  10.1.101.14            95   44d3.ca28.6cc4  ARPA   Vlan101 pv 103
Internet  10.1.101.15            95   44d3.ca28.6cc5  ARPA   Vlan101 pv 104
```

```
Internet  10.1.201.1              -   7c21.0dbd.954c  ARPA   Vlan201
Internet  172.16.254.4            -   7c21.0dbd.9548  ARPA   Vlan901

Leaf-02#
```

The following example shows the output for the **show mac address-table vlan** *vlan-id* command on VTEP 2:

```
Leaf-02# show mac address-table vlan 101
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
 101    10b3.d56a.8fc1    STATIC      Vl101
 101    44d3.ca28.6cc3    DYNAMIC pv  Gi1/0/11
 101    44d3.ca28.6cc4    DYNAMIC pv  Gi1/0/12
 101    44d3.ca28.6cc5    DYNAMIC pv  Gi1/0/13
 101    7c21.0dbd.9541    STATIC      Vl101
Total Mac Addresses for this criterion: 5

Leaf-02#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 2:

```
Leaf-02# show l2vpn evpn peers vxlan

Interface VNI      Peer-IP                 Num routes eVNI     UP time
--------- -------- ----------------------  ---------- -------- --------
nve1      10101    172.16.254.3            8          10101    01:34:10
nve1      10102    172.16.254.3            1          10102    01:34:10
nve1      10103    172.16.254.3            1          10103    01:34:10
nve1      10104    172.16.254.3            1          10104    00:02:13

Leaf-02#
```

The following example shows the output for the **show nve peer** command on VTEP 2:

```
Leaf-02# show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/4 01:34:10
nve1       50901    L3CP 172.16.254.5     7c21.0dbd.2748 50901      UP  A/M/4 01:34:10
nve1       10101    L2CP 172.16.254.3     8              10101      UP   N/A  01:34:10
nve1       10102    L2CP 172.16.254.3     1              10102      UP   N/A  01:34:10
nve1       10103    L2CP 172.16.254.3     1              10103      UP   N/A  01:34:10
nve1       10104    L2CP 172.16.254.3     1              10104      UP   N/A  00:02:13

Leaf-02#
```

The following example shows the output for the **show l2vpn evpn mac local** command on VTEP 2:

```
Leaf-02# show l2vpn evpn mac local
MAC Address    EVI   VLAN  ESI                       Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- ---------------
44d3.ca28.6cc3 101   101   0000.0000.0000.0000.0000 0          Gi1/0/11
44d3.ca28.6cc4 101   101   0000.0000.0000.0000.0000 0          Gi1/0/12
44d3.ca28.6cc5 101   101   0000.0000.0000.0000.0000 0          Gi1/0/13
44d3.ca28.6cc3 102   102   0000.0000.0000.0000.0000 0          Gi1/0/11
```

```
44d3.ca28.6cc4 103   103   0000.0000.0000.0000.0000 0          Gi1/0/12
44d3.ca28.6cc5 104   104   0000.0000.0000.0000.0000 0          Gi1/0/13

Leaf-02#
```

The following example shows the output for the **show l2vpn evpn mac remote** command on VTEP 2:

```
Leaf-02# show l2vpn evpn mac remote
MAC Address    EVI   VLAN  ESI                       Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------ ---------- --------------
f4cf.e243.34c2 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c3 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c4 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c2 102   102   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c3 103   103   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c4 104   104   0000.0000.0000.0000.0000 0          172.16.254.3
44d3.ca28.6cc6 201   201   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc7 201   201   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc8 201   201   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc6 202   202   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc7 203   203   0000.0000.0000.0000.0000 0          172.16.254.5
44d3.ca28.6cc8 204   204   0000.0000.0000.0000.0000 0          172.16.254.5

Leaf-02#
```

The following example shows the output for the **show l2route evpn mac ip** command on VTEP 2:

```
Leaf-02# show l2route evpn mac ip
  EVI       ETag  Prod    Mac Address         Host IP                Next Hop(s)
----- ---------- ----- -------------- --------------- --------------------------
  101          0   BGP 10b3.d56a.8fc1      10.1.101.1       V:10101 172.16.254.3
  101          0 L2VPN 44d3.ca28.6cc3     10.1.101.13               Gi1/0/11
  101          0 L2VPN 44d3.ca28.6cc4     10.1.101.14               Gi1/0/12
  101          0 L2VPN 44d3.ca28.6cc5     10.1.101.15               Gi1/0/13
  101          0 L2VPN 7c21.0dbd.9541      10.1.101.1                    Vl101:0
  101          0   BGP f4cf.e243.34c2      10.1.101.3       V:10101 172.16.254.3
  101          0   BGP f4cf.e243.34c3      10.1.101.4       V:10101 172.16.254.3
  101          0   BGP f4cf.e243.34c4      10.1.101.5       V:10101 172.16.254.3
  201          0   BGP 44d3.ca28.6cc6      10.1.102.3       V:10201 172.16.254.5
  201          0   BGP 44d3.ca28.6cc7      10.1.102.4       V:10201 172.16.254.5
  201          0   BGP 44d3.ca28.6cc8      10.1.102.5       V:10201 172.16.254.5
  201          0   BGP 7c21.0dbd.274c      10.1.201.1       V:10201 172.16.254.5

Leaf-02#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 2:

```
Leaf-02# show bgp l2vpn evpn
BGP table version is 65, local router ID is 172.16.255.4
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.255.3:101
 * i  [2][172.16.255.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
```

```
 *>i                    172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:102
 * i  [2][172.16.255.3:102][0][48][F4CFE24334C2][0][*]/20
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:103
 * i  [2][172.16.255.3:103][0][48][F4CFE24334C3][0][*]/20
                       172.16.254.3             0    100      0 ?
 *>i                    172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:104
 *>i  [2][172.16.255.3:104][0][48][F4CFE24334C4][0][*]/20
                       172.16.254.3             0    100      0 ?
 * i                    172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:101
 *>i  [2][172.16.255.4:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                       172.16.254.3             0    100      0 ?
 *>   [2][172.16.255.4:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                       ::                                    32768 ?
 *>   [2][172.16.255.4:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                       ::                                    32768 ?
 *>   [2][172.16.255.4:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                       ::                                    32768 ?
 *>   [2][172.16.255.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                       ::                                    32768 ?
 *>i  [2][172.16.255.4:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                       172.16.254.3             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
     Network          Next Hop            Metric LocPrf Weight Path
                       172.16.254.3             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                       172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:102
 *>   [2][172.16.255.4:102][0][48][44D3CA286CC3][0][*]/20
                       ::                                    32768 ?
 *>i  [2][172.16.255.4:102][0][48][F4CFE24334C2][0][*]/20
                       172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:103
 *>   [2][172.16.255.4:103][0][48][44D3CA286CC4][0][*]/20
                       ::                                    32768 ?
 *>i  [2][172.16.255.4:103][0][48][F4CFE24334C3][0][*]/20
                       172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:104
 *>   [2][172.16.255.4:104][0][48][44D3CA286CC5][0][*]/20
                       ::                                    32768 ?
 *>i  [2][172.16.255.4:104][0][48][F4CFE24334C4][0][*]/20
                       172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:201
 *>i  [2][172.16.255.4:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                       172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.4:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                       172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.4:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                       172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.4:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                       172.16.254.5             0    100      0 ?
```

```
Route Distinguisher: 172.16.255.4:202
 *>i  [2][172.16.255.4:202][0][48][44D3CA286CC6][0][*]/20
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:203
 *>i  [2][172.16.255.4:203][0][48][44D3CA286CC7][0][*]/20
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:204
 *>i  [2][172.16.255.4:204][0][48][44D3CA286CC8][0][*]/20
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:201
 *>i  [2][172.16.255.5:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.5:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.5:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?
 *>i  [2][172.16.255.5:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                      172.16.254.5             0    100      0 ?
     Network          Next Hop            Metric LocPrf Weight Path
 * i                    172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:202
 *>i  [2][172.16.255.5:202][0][48][44D3CA286CC6][0][*]/20
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:203
 *>i  [2][172.16.255.5:203][0][48][44D3CA286CC7][0][*]/20
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:204
 *>i  [2][172.16.255.5:204][0][48][44D3CA286CC8][0][*]/20
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][0][24][10.1.101.0]/17
                      172.16.254.3             0    100      0 ?
 * i                    172.16.254.3             0    100      0 ?
 *>                     0.0.0.0                  0         32768 ?
 *>i  [5][1:1][0][24][10.1.201.0]/17
                      172.16.254.5             0    100      0 ?
 * i                    172.16.254.5             0    100      0 ?

Leaf-02#
```

Outputs to Verify the Configuration on VTEP 3

The following example shows the output for the **show vlan private-vlan** command on VTEP 3:

```
Leaf-03# show vlan private-vlan

Primary Secondary Type              Ports
------- --------- ----------------- ------------------------------------------
101     102       community
101     103       community
101     104       isolated
201     202       community         Gi1/0/16
201     203       community         Gi1/0/17
201     204       isolated          Gi1/0/18

Leaf-03#
```

The following example shows the output for the **show ip arp vrf green** command on VTEP 3:

```
Leaf-03# show ip arp vrf green
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.1.101.1               -   7c21.0dbd.2741  ARPA   Vlan101
Internet  10.1.201.1               -   7c21.0dbd.274c  ARPA   Vlan201
Internet  172.16.254.5             -   7c21.0dbd.2748  ARPA   Vlan901

Leaf-03#
```

The following example shows the output for the **show mac address-table vlan** *vlan-id* command on VTEP 3:

```
Leaf-03# show mac address-table vlan 101
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
 101    7c21.0dbd.2741    STATIC      Vl101
Total Mac Addresses for this criterion: 1

Leaf-03#
```

The following example shows the output for the **show l2vpn evpn peers vxlan** command on VTEP 3:

```
Leaf-03# show l2vpn evpn peers vxlan

Leaf-03#
```

The following example shows the output for the **show nve peer** command on VTEP 3:

```
Leaf-03# show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/4 01:34:51
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/4 01:34:51

Leaf-03#
```

The following example shows the output for the **show l2vpn evpn mac local** command on VTEP 3:

```
Leaf-03# show l2vpn evpn mac local
MAC Address    EVI   VLAN  ESI                       Ether Tag  Next Hop(s)
-------------- ----- ----- ------------------------- ---------- ---------------
44d3.ca28.6cc6 201   201   0000.0000.0000.0000.0000 0          Gi1/0/16
44d3.ca28.6cc7 201   201   0000.0000.0000.0000.0000 0          Gi1/0/17
44d3.ca28.6cc8 201   201   0000.0000.0000.0000.0000 0          Gi1/0/18
44d3.ca28.6cc6 202   202   0000.0000.0000.0000.0000 0          Gi1/0/16
44d3.ca28.6cc7 203   203   0000.0000.0000.0000.0000 0          Gi1/0/17
44d3.ca28.6cc8 204   204   0000.0000.0000.0000.0000 0          Gi1/0/18

Leaf-03#
```

The following example shows the output for the **show l2vpn evpn mac remote** command on VTEP 3:

```
Leaf-03# show l2vpn evpn mac remote
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop(s)
-------------- ----- ----- ---------------------------- ---------- --------------
44d3.ca28.6cc3 101   101   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc4 101   101   0000.0000.0000.0000.0000 0          172.16.254.4
44d3.ca28.6cc5 101   101   0000.0000.0000.0000.0000 0          172.16.254.4
f4cf.e243.34c2 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c3 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
f4cf.e243.34c4 101   101   0000.0000.0000.0000.0000 0          172.16.254.3
44d3.ca28.6cc3 102   102   0000.0000.0000.0000.0000 0          172.16.254.4
f4cf.e243.34c2 102   102   0000.0000.0000.0000.0000 0          172.16.254.3
44d3.ca28.6cc4 103   103   0000.0000.0000.0000.0000 0          172.16.254.4
f4cf.e243.34c3 103   103   0000.0000.0000.0000.0000 0          172.16.254.3
44d3.ca28.6cc5 104   104   0000.0000.0000.0000.0000 0          172.16.254.4
f4cf.e243.34c4 104   104   0000.0000.0000.0000.0000 0          172.16.254.3

Leaf-03#
```

The following example shows the output for the **show l2route evpn mac ip** command on VTEP 3:

```
Leaf-03# show l2route evpn mac ip
  EVI       ETag  Prod    Mac Address          Host IP                 Next Hop(s)
----- ---------- ----- -------------- --------------- -------------------------
  101          0   BGP 10b3.d56a.8fc1      10.1.101.1       V:10101 172.16.254.3
  101          0   BGP 44d3.ca28.6cc3     10.1.101.13       V:10101 172.16.254.4
  101          0   BGP 44d3.ca28.6cc4     10.1.101.14       V:10101 172.16.254.4
  101          0   BGP 44d3.ca28.6cc5     10.1.101.15       V:10101 172.16.254.4
  101          0   BGP 7c21.0dbd.9541      10.1.101.1       V:10101 172.16.254.4
  101          0   BGP f4cf.e243.34c2      10.1.101.3       V:10101 172.16.254.3
  101          0   BGP f4cf.e243.34c3      10.1.101.4       V:10101 172.16.254.3
  101          0   BGP f4cf.e243.34c4      10.1.101.5       V:10101 172.16.254.3
  201          0 L2VPN 44d3.ca28.6cc6      10.1.102.3               Gi1/0/16
  201          0 L2VPN 44d3.ca28.6cc7      10.1.102.4               Gi1/0/17
  201          0 L2VPN 44d3.ca28.6cc8      10.1.102.5               Gi1/0/18
  201          0 L2VPN 7c21.0dbd.274c      10.1.201.1                    Vl201:0

Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn** command on VTEP 3:

```
Leaf-03# show bgp l2vpn evpn
BGP table version is 82, local router ID is 172.16.255.5
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.255.3:101
 * i  [2][172.16.255.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                      172.16.254.3             0    100      0 ?
```

```
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:102
 * i  [2][172.16.255.3:102][0][48][F4CFE24334C2][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:103
 * i  [2][172.16.255.3:103][0][48][F4CFE24334C3][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:104
 *>i  [2][172.16.255.3:104][0][48][F4CFE24334C4][0][*]/20
                      172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:101
 *>i  [2][172.16.255.4:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:102
     Network          Next Hop            Metric LocPrf Weight Path
 *>i  [2][172.16.255.4:102][0][48][44D3CA286CC3][0][*]/20
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:103
 *>i  [2][172.16.255.4:103][0][48][44D3CA286CC4][0][*]/20
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:104
 * i  [2][172.16.255.4:104][0][48][44D3CA286CC5][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:101
 *>i  [2][172.16.255.5:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.255.5:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                      172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.5:102
 *>i  [2][172.16.255.5:102][0][48][44D3CA286CC3][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.5:102][0][48][F4CFE24334C2][0][*]/20
                      172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.5:103
 *>i  [2][172.16.255.5:103][0][48][44D3CA286CC4][0][*]/20
                      172.16.254.4             0    100      0 ?
```

```
 *>i  [2][172.16.255.5:103][0][48][F4CFE24334C3][0][*]/20
                      172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.5:104
 *>i  [2][172.16.255.5:104][0][48][44D3CA286CC5][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.255.5:104][0][48][F4CFE24334C4][0][*]/20
                      172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.5:201
 *>   [2][172.16.255.5:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                      ::                                    32768 ?
 *>   [2][172.16.255.5:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                      ::                                    32768 ?
     Network          Next Hop            Metric LocPrf Weight Path
 *>   [2][172.16.255.5:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                      ::                                    32768 ?
 *>   [2][172.16.255.5:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                      ::                                    32768 ?
Route Distinguisher: 172.16.255.5:202
 *>   [2][172.16.255.5:202][0][48][44D3CA286CC6][0][*]/20
                      ::                                    32768 ?
Route Distinguisher: 172.16.255.5:203
 *>   [2][172.16.255.5:203][0][48][44D3CA286CC7][0][*]/20
                      ::                                    32768 ?
Route Distinguisher: 172.16.255.5:204
 *>   [2][172.16.255.5:204][0][48][44D3CA286CC8][0][*]/20
                      ::                                    32768 ?
Route Distinguisher: 1:1 (default for vrf green)
 * i  [5][1:1][0][24][10.1.101.0]/17
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 *>   [5][1:1][0][24][10.1.201.0]/17
                      0.0.0.0                  0         32768 ?

Leaf-03#
```

Outputs to Verify the Configuration on Spine Switch 1

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 113, main routing table version 113
23 network entries using 8832 bytes of memory
47 path entries using 10528 bytes of memory
15/14 BGP path/bestpath attribute entries using 4440 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
15 BGP extended community entries using 720 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 24640 total bytes of memory
BGP activity 47/24 prefixes, 107/60 paths, scan interval 60 secs
25 networks peaked at 13:03:03 Feb 19 2021 UTC (03:26:23.575 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.2    4        65001     259     261      113    0    0 03:27:45        23
172.16.255.3    4        65001     240     250      113    0    0 03:27:49         8
172.16.255.4    4        65001     238     258      113    0    0 03:27:25         8
172.16.255.5    4        65001     236     258      113    0    0 03:27:19         8

Spine-01#
```

The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 1:

```
Spine-01# show bgp l2vpn evpn
BGP table version is 113, local router ID is 172.16.255.1
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, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.255.3:101
 * i  [2][172.16.255.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:102
 * i  [2][172.16.255.3:102][0][48][F4CFE24334C2][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:103
 * i  [2][172.16.255.3:103][0][48][F4CFE24334C3][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:104
 * i  [2][172.16.255.3:104][0][48][F4CFE24334C4][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:101
 * i  [2][172.16.255.4:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.255.4:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.255.4:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.255.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:102
     Network          Next Hop            Metric LocPrf Weight Path
 * i  [2][172.16.255.4:102][0][48][44D3CA286CC3][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:103
 * i  [2][172.16.255.4:103][0][48][44D3CA286CC4][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:104
 * i  [2][172.16.255.4:104][0][48][44D3CA286CC5][0][*]/20
                      172.16.254.4             0    100      0 ?
```

```
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:201
 * i  [2][172.16.255.5:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.255.5:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.255.5:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.255.5:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:202
 * i  [2][172.16.255.5:202][0][48][44D3CA286CC6][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:203
 * i  [2][172.16.255.5:203][0][48][44D3CA286CC7][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:204
 * i  [2][172.16.255.5:204][0][48][44D3CA286CC8][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][24][10.1.201.0]/17
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?

Spine-01#
```

Outputs to Verify the Configuration on Spine Switch 2

The following example shows the output for the **show bgp l2vpn evpn summary** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 113, main routing table version 113
23 network entries using 8832 bytes of memory
47 path entries using 10528 bytes of memory
15/14 BGP path/bestpath attribute entries using 4440 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
15 BGP extended community entries using 720 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 24640 total bytes of memory
BGP activity 46/23 prefixes, 107/60 paths, scan interval 60 secs
25 networks peaked at 13:03:07 Feb 19 2021 UTC (03:27:53.810 ago)

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
172.16.255.1    4        65001     263     261      113    0    0 03:29:16        23
172.16.255.3    4        65001     243     251      113    0    0 03:29:17         8
172.16.255.4    4        65001     240     259      113    0    0 03:28:48         8
172.16.255.5    4        65001     240     257      113    0    0 03:28:45         8
```

```
Spine-02#
```

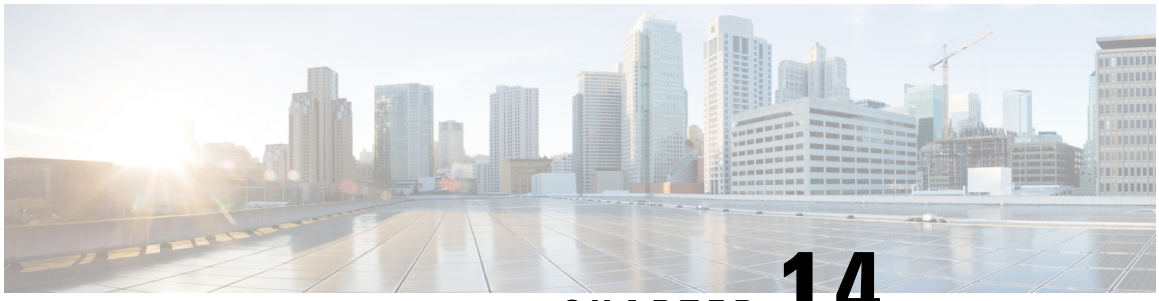
The following example shows the output for the **show bgp l2vpn evpn** command on Spine Switch 2:

```
Spine-02# show bgp l2vpn evpn
BGP table version is 113, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.255.3:101
 * i  [2][172.16.255.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C2][32][10.1.101.3]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C3][32][10.1.101.4]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.255.3:101][0][48][F4CFE24334C4][32][10.1.101.5]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:102
 * i  [2][172.16.255.3:102][0][48][F4CFE24334C2][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:103
 * i  [2][172.16.255.3:103][0][48][F4CFE24334C3][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.3:104
 * i  [2][172.16.255.3:104][0][48][F4CFE24334C4][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.255.4:101
 * i  [2][172.16.255.4:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.255.4:101][0][48][44D3CA286CC4][32][10.1.101.14]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.255.4:101][0][48][44D3CA286CC5][32][10.1.101.15]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
 * i  [2][172.16.255.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:102
     Network          Next Hop            Metric LocPrf Weight Path
 * i  [2][172.16.255.4:102][0][48][44D3CA286CC3][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.4:103
 * i  [2][172.16.255.4:103][0][48][44D3CA286CC4][0][*]/20
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
```

```
Route Distinguisher: 172.16.255.4:104
 * i  [2][172.16.255.4:104][0][48][44D3CA286CC5][0][*]/20
                     172.16.254.4             0    100      0 ?
 *>i                  172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:201
 * i  [2][172.16.255.5:201][0][48][44D3CA286CC6][32][10.1.102.3]/24
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
 * i  [2][172.16.255.5:201][0][48][44D3CA286CC7][32][10.1.102.4]/24
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
 * i  [2][172.16.255.5:201][0][48][44D3CA286CC8][32][10.1.102.5]/24
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
 * i  [2][172.16.255.5:201][0][48][7C210DBD274C][32][10.1.201.1]/24
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:202
 * i  [2][172.16.255.5:202][0][48][44D3CA286CC6][0][*]/20
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:203
 * i  [2][172.16.255.5:203][0][48][44D3CA286CC7][0][*]/20
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.5:204
 * i  [2][172.16.255.5:204][0][48][44D3CA286CC8][0][*]/20
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                     172.16.254.4             0    100      0 ?
 * i                  172.16.254.3             0    100      0 ?
 *>i                  172.16.254.3             0    100      0 ?
 * i  [5][1:1][0][24][10.1.201.0]/17
                     172.16.254.5             0    100      0 ?
 *>i                  172.16.254.5             0    100      0 ?

Spine-02#
```



CHAPTER 14

Configuring EVPN Microsegmentation

A secure BGP EVPN VXLAN fabric integrates Cisco TrustSec to provide microsegmentation and group-based policy enforcement.

- Restrictions for EVPN Microsegmentation, on page 633
- Information About EVPN Microsegmentation, on page 633
- How to Configure EVPN Microsegmentation, on page 636
- Deployment of EVPN Microsegmentation, on page 636
- Configuration Example of EVPN Microsegmentation, on page 640

Restrictions for EVPN Microsegmentation

Restrictions for EVPN Microsegmentation

- VXLAN-GPO encapsulation is not supported for VXLANv6. It is supported only with IPv4 underlay.
- If the host-facing interface is a trunk interface, you must configure device tracking on that interface.
- Overlay multicast traffic is not subject to SGACL enforcement.

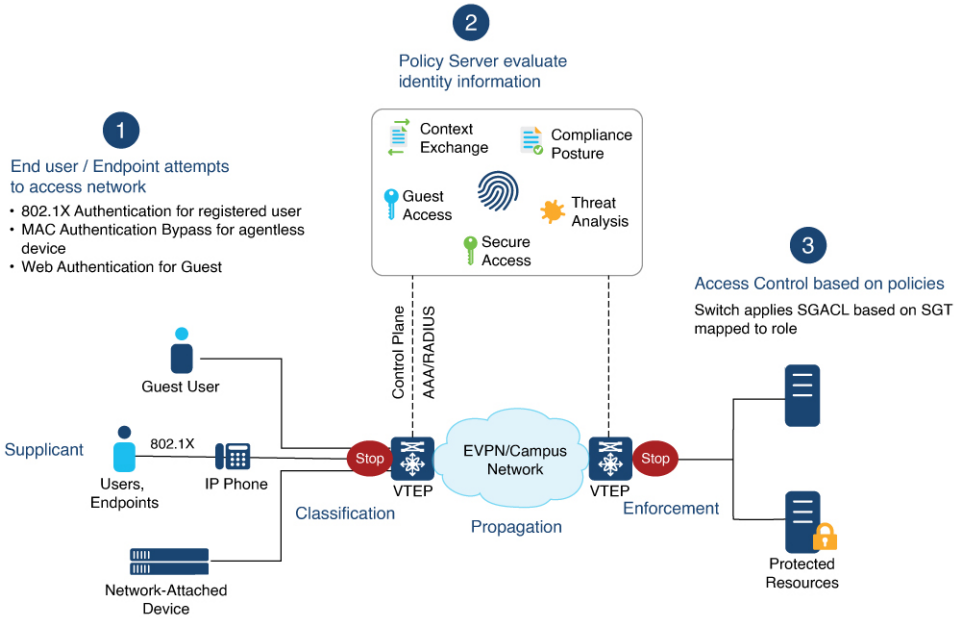
Information About EVPN Microsegmentation

In an enterprise network, users, devices, and applications, all utilize the network to access resources. But the access needs of a user may be different from that of a device in a network. Microsegmentation addresses the need for this isolation of network access. Using Security Group-Based Access Control, the endpoints within the overlay network can be permitted access to specific resources and denied access to others, based on their group membership.

BGP EVPN VXLAN integrates Cisco TrustSec to provide microsegmentation and end-to-end access control with propagation of the security group tag (SGT). Using security group access control lists (SGACLs), a network administrator can control the operations that users can perform based on their security group assignments and destination resources in a VXLAN campus fabric.

For more information on Cisco TrustSec, refer to “Cisco TrustSec Overview” in the *Cisco TrustSec Configuration Guide*.

Figure 59: Secure Fabric: BGP EVPN VXLAN Fabric Microsegmentation



Cisco TrustSec solution encompasses many aspects as discussed in the Cisco TrustSec Overview module. The three fundamental components of CiscoTrustSec are: Classification, Propagation, and Enforcement.

Classification

When users or devices (endpoints) connect to a network, the network assigns it a specific security group tag (SGT). This process is called Classification and can be either static or dynamic assignment. Static assignment is done by associating the SGT with an IP, VLAN, or port-profile. Dynamic assignment of the SGT tag is based on the results of authentication of the endpoints and downloaded as an authorization option from Cisco Identity Services Engine (ISE).

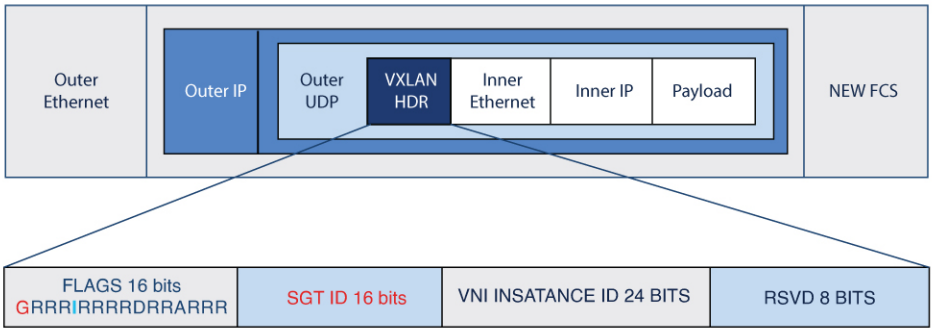
For information about ISE, refer to Cisco Identity Services Engine.

Propagation

An SGT is a metadata value that is transmitted in the header of the VXLAN-encapsulated packets. It is a unique 16-bit tag that represents the privilege of the source endpoint.

After the user traffic is classified, the SGT is propagated from the node at which classification took place, to where the enforcement action is invoked. This process is called propagation. Propagation in a BGP EVPN VXLAN network occurs through the VXLAN Group-based Policy Option (VXLAN-GPO), as defined in VXLAN Group Policy Option - draft-smith-vxlan-group-policy-0. SGT propagation through the EVPN VXLAN fabric is described in the later sections of this document.

A typical VXLAN header that contains an SGT tag is as shown:



Flag	Description
G Bit (bit 0)	Group-Based Policy Extension bit: Set to 1 when SGT is carried in the Group Policy ID field.
D bit (bit 9)	Don't Learn bit: Set to 1 to indicate that the egress VTEP must not learn the source address of the encapsulated frame. D bit is set to 0.
A bit (bit 12)	Policy Applied bit: A bit can be set to 0 (treat as reserved).
Group Policy ID	The Group Policy ID is a 16-bit identifier that identifies the source group of the packet encapsulated within VXLAN. Set to SGT (also called EPG Class).
VXLAN Network Identifier (VNI)	VNI is a 24-bit identifier that identifies the virtual network of the communicating hosts. It is set to the egress VNI associated with the destination subnet as determined by routing or bridging.

Enforcement

The enforcement of the policy based on the SGT occurs at an egress device like a firewall, router, or a switch. The enforcement device takes the source SGT and looks it up against the destination SGT to determine if the traffic should be allowed or denied. Security group access control list (SGACL) is one of the policy enforcement methods that provides state-less access control mechanism based on the security association or SGT value.

Security Group Access Control Lists

SGACLs define access control policies based on device identities instead of IP addresses as in traditional ACLs. Hence the network devices are free to move throughout the network and change IP addresses. If the roles and the permissions remain the same, changes to the network topology do not change the security policy. When an endpoint is connected to the device, you simply assign the endpoint to an appropriate security group and the endpoint immediately receives the permissions of that group.Using role-based permissions greatly reduces the size of ACLs and simplifies their maintenance. With Cisco TrustSec, the number of access control entries (ACEs) configured is determined by the number of permissions specified, resulting in a much smaller number of ACEs than in a traditional IP network. The use of SGACLs typically results in a more efficient use of TCAM resources compared with traditional ACLs. For more information on the SGACL policies and

enforcement, refer to “Configuring Security Group ACL Policies” in the *Cisco TrustSec Configuration Guide* for the relevant platform.

How to Configure EVPN Microsegmentation

Procedure

Step 1 **Classify**: Authenticate an endpoint and authorize it’s the network access by assigning an SGT tag.

a) For configuration procedures to authenticate an endpoint at ingress, refer to the “Configuring Endpoint Admission Control” chapter in the *Cisco TrustSec Configuration Guide*.

b) For configuration procedures to assign an SGT to a device or a user at ingress, refer to the "Configuring SGT Exchange Protocol" chapter in the *Cisco TrustSec Configuration Guide*.

Step 2 **Propagate**: Enable the propagation of the SGT information using the **group-based-policy** command on the NVE interface.

Example:

```
interface nve10
 no ip address
 source-interface Loopback0
 host-reachability protocol bgp
 group-based-policy
```

After the VXLAN-GPO encapsulation is enabled, it is applicable to all network overlay segments.

Step 3 **Enforce**: Define the SGACL policies and enforce them through egress filtering of the packets.

For the procedure to configure SGACL policy, refer to "Configure Security Group ACL Polices” chapter of the *Cisco TrustSec Configuration Guide*.

Deployment of EVPN Microsegmentation

The EVPN VXLAN fabric propagates the SGT tag (and thus security policy) in the following deployment scenarios:

- **EVPN Layer 2 Virtual Network Instance (L2VNI)**

  An EVPN VXLAN Layer 2 overlay network allows host devices in the same subnet to send bridged traffic to each other. The network forwards the bridged traffic using a Layer 2 virtual network instance (VNI). Layer 2 overlay network traffic supports the SGT propagation.

  To configure a Layer 2 overlay, refer to How to Configure EVPN VXLAN Layer 2 Overlay Network, on page 16.

- **EVPN Layer 3 Virtual Network Instance (L3VNI)**

An EVPN VXLAN Layer 3 overlay network allows host devices in different Layer 2 networks to send Layer 3 or routed traffic to each other. The network forwards the routed traffic using a Layer 3 virtual network instance (VNI) and an IP VRF.

For information about how to configure a Layer 3 overlay network, refer to How to Configure EVPN VXLAN Layer 3 Overlay Network.

- **EVPN VXLAN Distributed Anycast Gateway**

  EVPN VXLAN integrated routing and bridging (IRB) allows the VTEPs or leaf switches in an EVPN VXLAN network to perform both bridging and routing. The VTEPs in the network forward traffic to each other through the VXLAN gateways.

  Distributed anycast gateway enables the use of the same gateway IP and MAC address across all the VTEPs in an EVPN VXLAN network. This ensures that every VTEP functions as the default gateway for the workloads directly connected to it.

  To configure a Distributed Anycast Gateway, refer to How to Configure EVPN VXLAN IRB using Anycast Distributed Gateway

- **EVPN VXLAN Centralized Anycast Gateway**

  A Centralized Anycast Gateway (CGW) VTEP performs the Layer 3 gateway function for all the Layer 2 VNIs. All the other VTEPs in the network perform only bridging. The CGW VTEP acts as the Layer 3 gateway and performs routing for the intersubnet VXLAN traffic.

  To configure a CGW, refer to How to Configure EVPN VXLAN IRB using CGW.

- **Multi-Homing**

  Multi-homing allows you to connect a host or Layer 2 switch to more than one VTEP in the EVPN VXLAN network. This connection provides redundancy and allows network optimization. Redundancy in the connection with the VTEPs ensures that there's no traffic disruption when there's a network failure.

  To configure Multi-homing, refer to How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric.

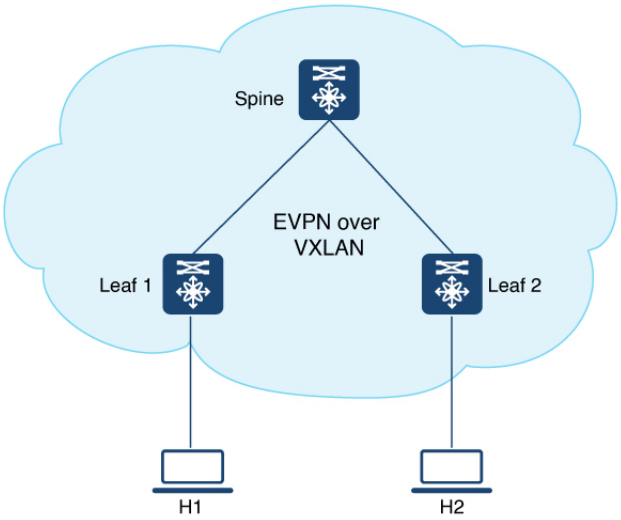
In each of the deployments, the leaf node is located either at the access layer or the distribution layer.

Leaf Node at Access or Distribution Layer

Leaf Node at Access Layer

The CTS solution is enabled on the leaf node to provide classification of packets and secure propagation and enforcement of policies. When a host/endpoint connects to a Leaf node that is located at the access layer:

- The leaf node first authenticates the endpoint using 802.1X protocol.
- On successful authentication, the endpoint is authorized to send and receive traffic. Traffic from the endpoint is classified and tagged with an SGT value.



Configuration of the Leaf Node

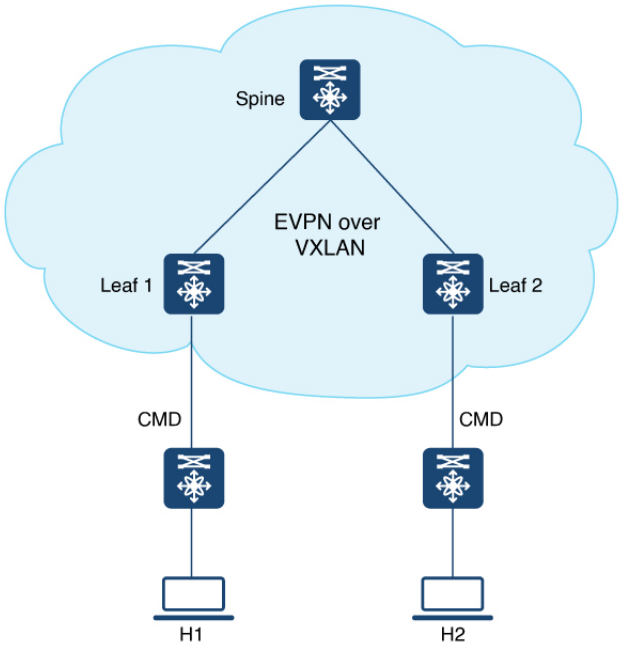
The leaf node is configured to provide the functionalities:

- VXLAN-GPO encapsulation
- CTS classification, propagation, and enforcement

Leaf Node at the Distribution Layer

The CTS solution is enabled on the access switch. When an endpoint connects to the switch at the access layer:

- Endpoint is authenticated with the IEEE 802.1X method.
- On successful authentication, the traffic from the endpoint is classified and tagged with an SGT value.
- The packets originating from the endpoints are encapsulated in a Cisco Meta Data (CMD) format at the access layer and forwarded to the leaf node at the distribution layer (L2VNI). The CMD header is inserted in the frame header before being sent out of the access switch to the VTEP (leaf node) in the fabric.
- The Leaf node encapsulates the packet with a VXLAN header, preserving the CMD data and sends the packet to the destination.



Configuration of Leaf Node and Access Switch

The leaf node at the distribution layer is configured to provide VXLAN-GPO encapsulation of the packets at egress, and propagation of SGT tags.

The access switch is configured to provide CTS classification.

Device Tracking for a Trunk Interface

When you configure the group-based policy, you must explicitly attach a device-tracking policy to the trunk interface of the Layer 2 VTEP. This explicit policy attachment is required to keep the IP-to-SGT binding active.

The following sample configuration shows how to define a policy and attach it to a specific target, which is the trunk interface.

1. Define a new policy.

```
VTEP(config)# device-tracking policy IPDT_POLICY
VTEP(config-device-tracking)# tracking enable
VTEP(config-device-tracking)# security-level glean
VTEP(config-device-tracking)# exit
```

2. Verify the policy that is created.

```
VTEP# show device-tracking policy IPDT_POLICY
Device-tracking policy IPDT_POLICY configuration:
security-level glean
device-role node
gleaning from Neighbor Discovery
gleaning from DHCP6
gleaning from ARP
```

```
gleaning from DHCP4
NOT gleaning from protocol unknown
tracking enable
```

3. Configure the policy on a specific target (trunk interface).

```
VTEP(config)# interface gi1/0/1
VTEP(config-if)# switchport trunk allowed vlan 101,102,201,202
VTEP(config-if)# switchport mode trunk
VTEP(config-if)# device-tracking attach-policy IPDT_POLICY
VTEP(config-if)# exit
```

Configuration Example of EVPN Microsegmentation

Configuration of SGACL policies should be done primarily through the Policy Management function of the Cisco Secure Access Control Server (ACS) or the Cisco Identity Services Engine (ISE).

If you are not using AAA on a Cisco Secure ACS or a Cisco ISE to download the SGACL policy configuration, you can manually configure the SGACL mapping and policies.

The following is a snippet of a sample configuration of static SGT map on a Leaf node.

```
<snip: only relevant configuration shown>
cts role-based sgt-map 192.168.0.0/16 sgt 192
cts role-based sgt-map 192:168::/64 sgt 192
cts role-based sgt-map vrf green 10.2.201.10 sgt 101
cts role-based sgt-map vlan-list 101 sgt 101
cts role-based sgt-map vlan-list 102 sgt 102
<snip: only relevant configuration shown>
```

The following is a snippet of a sample configuration on the VTEPs for CTS enforcement.

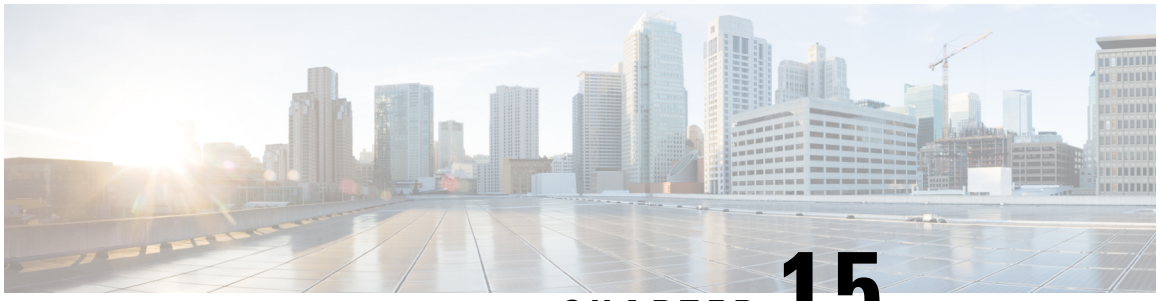
```
<snip: only relevant configuration shown>
cts role-based enforcement
cts role-based enforcement vlan-list 101
cts role-based enforcement logging-interval 5
cts role-based permissions default d_permit
cts role-based permissions from 0 to 0 m_icmp
cts role-based permissions from 31 to 31 m_permit
cts role-based permissions from 101 to 192 m_icmp m_permit
<snip: only relevant configuration shown>
```

The following is a snippet of a sample configuration on the NVE interface that enables propagation of SGT information.

```
<snip: only relevant configuration shown>
interface nve1
  group-based-policy
<snip: only relevant configuration shown>
```

The following is a snippet of a sample configuration to accept SGTs carried via CMD header.

```
<snip: only relevant configuration shown>
interface GigabitEthernet1/0/1
 switchport trunk allowed vlan 101,102,201,202
switchport mode trunk
 cts manual
  policy static sgt 2 trusted
<snip: only relevant configuration shown>
```



CHAPTER 15

Configuring Additional BGP EVPN Features

This module discusses some features of BGP EVPN such as Dynamic BGP Peering, Route Map support, and so on.

- Dynamic BGP Peering for EVPN, on page 641
- EVPN Route Map Support, on page 641
- BGP EVPN VRF Auto Route Distinguisher, on page 645
- BGP EVPN VRF Auto Route Target, on page 646

Dynamic BGP Peering for EVPN

Dynamic BGP Peering feature allows BGP peering to a group of remote neighbors that are defined by a range of IP addresses. Each range can be configured as a subnet IP address. BGP dynamic neighbors are configured using a range of IP addresses and BGP peer groups.

For information about configuring Dynamic BGP Peering, see *BGP Dynamic Neighbors*.

Configuration Example for Dynamic BGP Peering

Here is a sample configuration for dynamic BGP peering for L2VPN EVPN Family:

```
router bgp 10
 bgp log-neighbor-changes
 bgp listen range 10.10.10.0/24 peer-group DN-GRP
 bgp listen limit 2
 neighbor DN-GRP peer-group
 neighbor DN-GRP remote-as 12 alternate-as 13
 !
 address-family l2vpn evpn
  neighbor DN-GRP activate
  neighbor DN-GRP send-community extended
 exit-address-family
 !
```

EVPN Route Map Support

The Leaf, Spine and Border nodes of a BGP EVPN fabric support route map for L2VPN EVPN address-family. A route map allows you to check for certain match conditions and set a value. With the route map support,

the BGP attributes and their values can be modified to customize the routing policy based on the requirement. The routing policy can be applied for both inbound and outbound EVPN routes.

EVPN routes contain many fields and carry attributes specific to EVPN. Using route maps, you can filter routes on the basis of these attributes. The following route-filtering and attribute-handling options are available for the routes belonging to the EVPN address family:

- Matching based on IP prefix:

  You can use IPv4 or IPv6 prefix or host route that is embedded in EVPN prefix field to filter EVPN Type 2 and Type 5 routes.

- Matching based on the EVPN route type:

  Seven types of network layer reachability information (NLRI) are available in EVPN and called the Route Types. You can filter the routes based on the EVPN route type. For Type 2, you can further filter the routes based on the MAC address or the MAC and IP addresses.

- Setting a value:

  A set clause is used in a route map to set the BGP attributes, community, and so on. In a route map operation, after all the match statements are successful, the set clauses are executed.

Configuring Route Maps

You can configure route maps to filter the EVPN address-family routes. A route map evaluates match clauses until a successful match occurs. After the match occurs, the route map stops evaluating match clauses and starts executing set clauses, in the order in which they were configured. If a successful match does not occur, the route map "falls through" and evaluates the next sequence number of the route map until all configured route map entries have been evaluated or a successful match occurs.

Matching based on IP prefix

Use this procedure to filter the BGP EVPN routes based on the value of the EVPN prefix field. You can use the IPv4 or IPv6 prefix or host route that is embedded in the EVPN prefix field to filter the EVPN Type 2, Type 3, and Type 5 routes.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 2	**route-map** *route-map-name* **Example:** `Device(config)# route-map rmap1`	Creates a route map and enters route-map configuration mode.
Step 3	**match ip address prefix-list** *prefix-list-name* **Example:**	Matches the BGP EVPN routes with the specified prefix-list.

	Command or Action	Purpose
	`Device(config-route-map)# match ip address prefix-list plist1`	
Step 4	**end** **Example:** `Device(config-router-af)# end`	Exits route-map configuration mode and enters privileged EXEC mode.

Example

The following snippet is a sample configuration that matches an EVPN route based on IP prefix.

```
ip prefix-list plist1 permit 172.16.1.0/24
route-map rmap1
  match ip address prefix-list plist1
  set local-preference 222
```

Matching Based on EVPN Route Type

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 2	**route-map** *route-map-name* **Example:** `Device(config)# route-map rmap1`	Creates a route map and enters route-map configuration mode.
Step 3	**match evpn route-type** { **1** | **2** | **2-mac-only** | **2-mac-ip** | **3** | **4** | **5** | **6** } **Example:** `Device(config-route-map)# match evpn route-type 6`	Matches the BGP EVPN routes with the specified route type.
Step 4	**end** **Example:** `Device(config-router-af)# end`	Exits route-map configuration mode and enters privileged EXEC mode.

Example

The following snippet is a sample configuration that matches an EVPN route based on the route type 2.

```
route-map rmap1 permit 10
 match evpn route-type 2
 set local-preference 222
```

Applying Route Maps

You can apply the route maps to filter the EVPN routes that are advertised to or received from a BGP peer and to set BGP attributes to specific routes.

After you've configured the route map with match and set clauses based on your requirements, use the following procedure to apply the route map at the inbound or outbound level.

Procedure

	Command or Action	Purpose
Step 1	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 2	**router bgp** *autonomous-system* **Example:** `Device(config)# router bgp 100`	Enables a BGP routing process, assigns it an AS number, and enters router configuration mode. The AS number can be from 1 to 65535, with 64512 to 65535 designated as private autonomous numbers.
Step 3	**neighbor** {*ip-address* | *peer-group-name* } **remote-as** *number* **Example:** `Device(config-router)# neighbor 10.1.1.2 remote-as 200`	Adds an entry to the BGP neighbor table specifying that the neighbor that is identified by the IP address belongs to the specified AS.
Step 4	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Specifies the L2VPN address family and enters address family configuration mode.
Step 5	**neighbor** *ip-address* **activate** **Example:** `Device(config-router-af)# neighbor 10.1.1.2 activate`	Enables the exchange of information with a BGP neighbor.
Step 6	**neighbor** *ip-address* **route-map** *route-map-name* {**in** | **out**} **Example:** `Device(config-router-af)# neighbor 10.1.1.2 route-map rmap1 in`	Applies a route map to incoming or outgoing routes.

	Command or Action	Purpose
Step 7	**end** **Example:** `Device(config-router-af)# end`	Exits address-family configuration mode and enters privileged EXEC mode.

Example

The following example shows how to configure a route map as an inbound neighbor policy.

```
ip prefix-list plist1 permit 10.1.1.0/24
 route-map rmap1 permit 10
 match ip address prefix-list plist1
 set community comm1

 router bgp 100
 neighbor 10.1.1.2 remote-as 200
 address-family l2vpn evpn
  neighbor 10.1.1.2 activate
  neighbor 10.1.1.2 route-map rmap1 in
!
```

BGP EVPN VRF Auto Route Distinguisher

You can now configure devices in a BGP EVPN VXLAN fabric to autogenerate the route distinguisher (RD) value. Autogenerated RD value is based on the Type 1 encoding format as described in IETF RFC 4364.

Prerequisites for Configuring Auto Route Distinguisher in a BGP EVPN VXLAN Fabric

You must explicitly configure the BGP router ID at the global level, using the **bgp router-id** command.

RD is autogenerated using four bytes of BGP router ID and two bytes of VRF ID.

Configuring Auto Route Distinguisher in a BGP EVPN VXLAN Fabric

To enable autogeneration of RD for a specific VRF, use the **rd-auto** command in the VRF configuration mode.

Example

```
VTEP(config)# vrf definition green
VTEP(config-vrf)# rd-auto
```

To enable autogeneration of RDs at a global level for all the VRFs in the fabric, use the **vrf rd-auto** command in the global configuration mode.

Example

```
VTEP(config)# vrf rd-auto
```

To override the global auto RD configuration and configure a static RD for a specific VRF, use the **rd-auto disable** command in the VRF configuration mode. This command disables the autogeneration of RD for this specific VRF and lets you configure a static RD value.

Example

```
VTEP(config)# vrf definition green
VTEP(config-vrf)# rd-auto disable
VTEP(config-vrf)# rd 103:2
```

You can disable the manual configuration of RD using the **no** form of the command.

The following example disables the static configuration of RD and enables autogeneration of RD.

Example

```
VTEP(config)# vrf definition green
VTEP(config-vrf)# no rd 103:2
VTEP(config-vrf)# no rd-auto disable
```

Note You cannot enable both auto RD and static RD for a given VRF.

In , the **show vrf** command is enhanced to indicate if an RD has been assigned automatically.

Example

```
VTEP# show vrf
Name              Default RD            Protocols  Interfaces
green             1.1.1.1:60000(auto)   ipv4,ipv6   Gi1/0/1
blue              2:2                   ipv4,ipv6   Lo12
```

BGP EVPN VRF Auto Route Target

A VRF can be associated with one or more route targets (RT). You can enable the devices in a BGP EVPN VXLAN fabric to autogenerate the route target value for a specified VRF. The autogenerated route target value is based on the Type 0 encoding format as described in IETF RFC 4364.

Note Both statically assigned RT and autogenerated RT can exist in a VRF.

Prerequisites for Configuring Auto Route Target in a BGP EVPN Fabric

You must explicitly configure the BGP router-id at the global level, using the **bgp router-id** command in the router configuration mode.

A route target value is constructed using Autonomous System Number (ASN) as the 2-byte administrative field and the Virtual Network Identifier (VNID) for the 4-byte numbering field. If the ASN is 4 bytes, it is aliased with a special purpose 2-byte ASN, 23456. This 2-byte ASN 23456 is registered by the IANA as AS_TRANS.

Configuring Auto Route Target in a BGP EVPN VXLAN Fabric

To enable autogeneration of route target for a VRF, use the following command in the VRF configuration mode.

vnid *vnid* [**evpn-instance** [**vni** {*vni-id* | **auto**} **core-vlan** *vlan-id*]]

The following table describes the different fields of the command syntax.

vnid	Specifies the Virtual Network Identifier (VNID). The VNID range is from 1 to 2147483647. This number is primarily used to auto generate the route targets.
evpn-instance	(Optional) Specify the **evpn-instance** keyword to autogenerate the stitching of Route Targets.
vni *vni-id*	(Optional) Specifies the NVE Layer 3 VNI number. The range is from 4096 to 16777215.
vni auto	(Optional) Auto generates the NVE Layer 3 VNI number using the *vnid* value.
core-vlan *vlan-id*	(Optional) Auto generates the core VLAN configuration using *vnid* and *vlan-id*. *vlan-id* ranges from 1 to 4094. It also auto generates the core SVI which uses the IP address of the source-interface of the NVE interface. **Note** • If you use the **vni** option, it is mandatory to provide the **core-vlan** option too. • When the core SVI is autogenerated, if the NVE source interface is not configured with an IP address, the corresponding Layer 3 traffic does not flow because of a lack of address on the core SVI.

Example

```
VTEP(config)# router bgp 150
VTEP(config-bgp)# bgp router-id interface Loopback0
VTEP(config-bgp)# exit

VTEP(config)# vrf definition blue
VTEP(config-vrf)# vnid 123 evpn-instance vni 35000 core-vlan 345
VTEP(config-vrf)# exit

VTEP(config)# address-family ipv4
VTEP(config-af-ipv4)# exit-address-family
VTEP(config)# address-family ipv6
VTEP(config-af-ipv6)# exit-address-family
```

The **show vrf detail** command is enhanced to indicate if an RT has been assigned automatically.

Example

```
VTEP# show vrf detail blue
VRF blue (VRF Id = 2); default RD 192.168.1.1:1(auto); default VPNID <not set>
  New CLI format, supports multiple address-families
  vnid: 123 evpn-instance vni 35000 core-vlan 345
  Flags: 0x180C
  Interfaces:
    Vl345                   Lo101                   Et1/1
    Vl4                     Vl15
Address family ipv4 unicast (Table ID = 0x2):
  Flags: 0x0
  Export VPN route-target communities
    RT:100:123              RT:65000:123 (auto)
  Import VPN route-target communities
    RT:100:123              RT:65000:123 (auto)
  Export VPN route-target stitching communities
    RT:65000:123 (auto)
  Import VPN route-target stitching communities
    RT:65000:123 (auto)
  No import route-map
  No global export route-map
  No export route-map
  VRF label distribution protocol: not configured
  VRF label allocation mode: per-prefix
Address family ipv6 unicast not active
Address family ipv4 multicast not active
Address family ipv6 multicast not active

VTEP1#
```

Use the **show derived-config** command to check the acquired configurations for core VLAN, core SVI and the NVE interface.

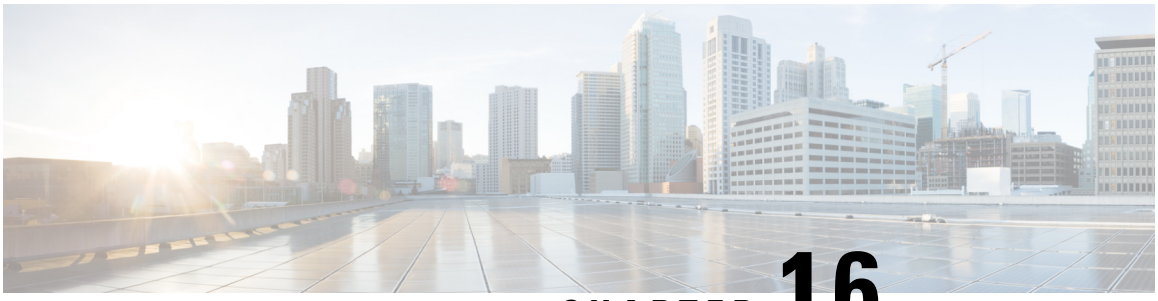
```
VTEP1# show derived-config interface Vlan345

Building configuration...
Derived configuration : 66 bytes
!
interface Vlan345
 vrf forwarding blue
 ip unnumbered Loopback1
```

```
 ipv6 unnumbered Loopback1
 no autostate
end

VTEP1# show derived-config interface nve1

Building configuration...
Derived configuration : 296 bytes
!
interface nve1
 source-interface Loopback1
 host-reachability protocol bgp
 member vni 35555 vrf black
 member vni 35000 vrf blue
 member vni 20011 ingress-replication
 member vni 20012 ingress-replication
 member vni 30000 vrf red
 member vni 30002 vrf green
 member vni 30004 vrf yellow
end
```



CHAPTER 16

Troubleshooting BGP EVPN VXLAN

- Troubleshooting Scenarios for BGP EVPN VXLAN, on page 651
- Troubleshooting Broadcast, Unkown Unicast, Multicast Traffic Forwarding, on page 652
- Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI, on page 656
- Troubleshooting Unicast Forwarding Between VTEPS in Different VLANs Through a Layer 3 VNI, on page 668
- Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network, on page 681
- Troubleshooting Tenant Routed Multicast, on page 688

Troubleshooting Scenarios for BGP EVPN VXLAN

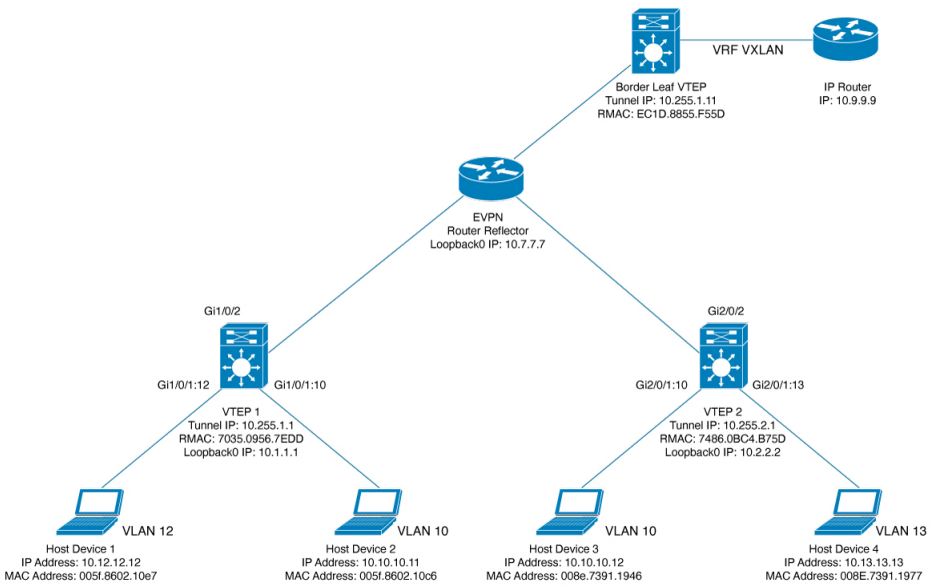
This document provides information about the various troubleshooting scenarios that are applicable to BGP EVPN VXLAN and how to troubleshoot each scenario.

In this troubleshooting document, comments have been added at the end of certain lines of the outputs of **show** commands. This has been done to highlight or explain a specific aspect of that line of output. If a comment begins in a new line, then it refers to the line of output that precedes the comment. The following notation has been used throughout the document to highlight the comments inside the outputs of **show** commands:

```
<<— Text highlighted in this format inside a command's output represents a comment.
    This is done for explanation purpose only and is not part of the command's output.
```

The following is a sample EVPN VXLAN topology with two access facing VTEPs (VTEP 1 and VTEP 2) and a border leaf VTEP connected in a VXLAN network through an EVPN route reflector. Each of the access facing VTEPs has two host devices connected to it and the border leaf VTEP is connected to an external IP network. All the troubleshooting scenarios in this document are explained using this topology.

Figure 60: EVPN VXLAN Topology



The following are the various troubleshooting scenarios that apply to BGP EVPN VXLAN for the topology illustrated in the figure above:

- **Scenario 1:** Troubleshooting Broadcast, Unknown Unicast, Multicast traffic Forwarding
- **Scenario 2:** Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI
- **Scenario 3:** Troubleshooting Unicast Forwarding Between VTEPS in Different VLANs Through a Layer 3 VNI
- **Scenario 4:** Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network

Troubleshooting Broadcast, Unkown Unicast, Multicast Traffic Forwarding

This scenario might occur when host device 2 attempts to learn the ARP for host device 3 in Troubleshooting Scenarios for BGP EVPN VXLAN, on page 651. Perform the checks listed in the following table before troubleshooting BUM traffic forwarding:

Table 57: Scenario 1: Broadcast, Unkown Unicast, Multicast traffic Forwarding

Check to be Performed	Steps to Follow
Is the packet of broadcast type?	Check if the packet is a broadcast packet, such as an ARP broadcast packet.

Check to be Performed	Steps to Follow
Are the hosts in the same subnet or in different subnets?	Perform any of the following steps: • Check the host device. • Check the SVI configuration on the VTEP.
Has the remote MAC address been learned for unknown unicast traffic?	Run the **show platform software fed switch active matm macTable vlan** *vlan-id* command in privileged EXEC mode on the local VTEP and check if the MAC address of the remote host device is displaed in the output. If not, you have not yet learned the remote host device and it needs to be resolved.

BUM traffic is forwarded by a VTEP into the VXLAN Core using multicast routing. In order to follow the path of an ARP broadcast packet, you need to identify the multicast group that needs to be used to send this traffic into the core and to the other VTEPs. BUM traffic first arrives at the local Layer 2 interface. The traffic is encapsulated here and sent out using the multicast group that is sourced from the VXLAN Loopback interface.

Note Underlay multicast needs to be fully configured before troubleshooting BUM traffic forwarding for EVPN VXLAN.

To troubleshoot EVPN VXLAN BUM traffic forwarding, follow these steps:

1. Determine the MAC Address of the Local Host Device and the Multicast Group Used for ARP Tunneling, on page 653
2. Set Up Embedded Capture Towards the Core-Facing Interface, on page 654
3. Ping the Remote Host Device, on page 654
4. Verify that an ARP Request Has Been Received and a Multicast Route Has Been Built, on page 654
5. Confirm the Presence of ARP Request Replies in Embedded Capture, on page 655
6. Verify that the Encapsulated ARP Request is Leaving in a Multicast Group to a VXLAN UDP Destination Port, on page 655
7. Verify that the ARP Reply from Core Interface is Encapsulated in Unicast to a VXLAN UDP Destination Port, on page 656

Determine the MAC Address of the Local Host Device and the Multicast Group Used for ARP Tunneling

The following examples show how to verify the MAC address of the local host device and the multicast group that is used for tunneling the ARP broadcast request:

```
VTEP-1# show mac address-table address 005f.8602.10c6
Mac Address Table
-------------------------------------------

Vlan Mac Address Type     Ports
```

```
---- ----------- -------- -----
10 005f.8602.10c6 DYNAMIC Gi1/0/1   <<— MAC address of 10.10.10.11 is learnt here

VTEP-1# show run int nve 1
interface nve1
 no ip address
 source-interface Loopback999
 host-reachability protocol bgp
 member vni 10001 mcast-group 239.10.10.10    <<— Group is mapped to the VNI under NVE

VTEP-1# show run | s vlan conf
vlan configuration 10
 member evpn-instance 10 vni 10001   <<— VNI mapped under VLAN 10

VTEP-1# show l2vpn evpn evi
EVI   VLAN  Ether Tag  L2 VNI    Multicast     Pseudoport
----- ----- ---------- --------- ------------- ------------------
10    10    0          10001     239.10.10.10  Gi1/0/1
   <<— EVPN instance 10 is mapped to VLAN 10 and VNI 10001
       (Using multicast group 239.10.10.10 for Broadcast ecap tunnel)
<...snip...>
```

Set Up Embedded Capture Towards the Core-Facing Interface

The following example shows how to set up embedded capture towards the core-facing interface:

Note On a production network, use this command with a filter.

```
VTEP-1# show monitor capture 1 parameter
monitor capture 1 interface GigabitEthernet1/0/1 BOTH
monitor capture 1 match any
monitor capture 1 buffer size 100
monitor capture 1 limit pps 1000
```

Ping the Remote Host Device

The following example shows how to ping the remote host device:

```
VTEP-1-HOST# ping 10.10.10.12     <<— sourced from Host machine 10.10.10.11
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.10.10.12, timeout is 2 seconds:
..!!!
```

Verify that an ARP Request Has Been Received and a Multicast Route Has Been Built

This step is to verify that there is multicast reachability between VTEPs using standard multicast validation. Underly multicast state is not permanent. If it is not in use, these S,G states will expire.

The following output confirms that an ARP request has been received and a multicast route has been built:

```
VTEP-1# show ip mroute 239.10.10.10 10.255.1.1
IP Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group, C - Connected,
```

```
L - Local, P - Pruned, R - RP-bit set, F - Register flag,
T - SPT-bit set, J - Join SPT, M - MSDP created entry, E - Extranet,
X - Proxy Join Timer Running, A - Candidate for MSDP Advertisement,
U - URD, I - Received Source Specific Host Report,
Z - Multicast Tunnel, z - MDT-data group sender,
Y - Joined MDT-data group, y - Sending to MDT-data group,
G - Received BGP C-Mroute, g - Sent BGP C-Mroute,
N - Received BGP Shared-Tree Prune, n - BGP C-Mroute suppressed,
Q - Received BGP S-A Route, q - Sent BGP S-A Route,
V - RD & Vector, v - Vector, p - PIM Joins on route,
x - VxLAN group, c - PFP-SA cache created entry
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(10.255.1.1, 239.10.10.10), 00:00:25/00:02:34, flags: FTx   <<— x flag set for VxLAN group
Incoming interface: Loopback999, RPF nbr 0.0.0.0    <<— Broadcast being encapsulated
                                                        into VXLAN tunnel IP
Outgoing interface list:
GigabitEthernet1/0/1, Forward/Sparse, 00:00:23/00:03:06
   <<— Sending towards core to VTEP-2
(10.255.1.4, 239.10.10.10), 3d18h/00:02:25, flags: JTx <<— BUM traffic from VTEP-2 (if the
                                                           ARP request was from VTEP-2)
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 10.1.1.6
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 3d18h/00:00:14    <<— Tunnel 0 is the VXLAN tunnel
                                                          used for decapsulation
```

Confirm the Presence of ARP Request Replies in Embedded Capture

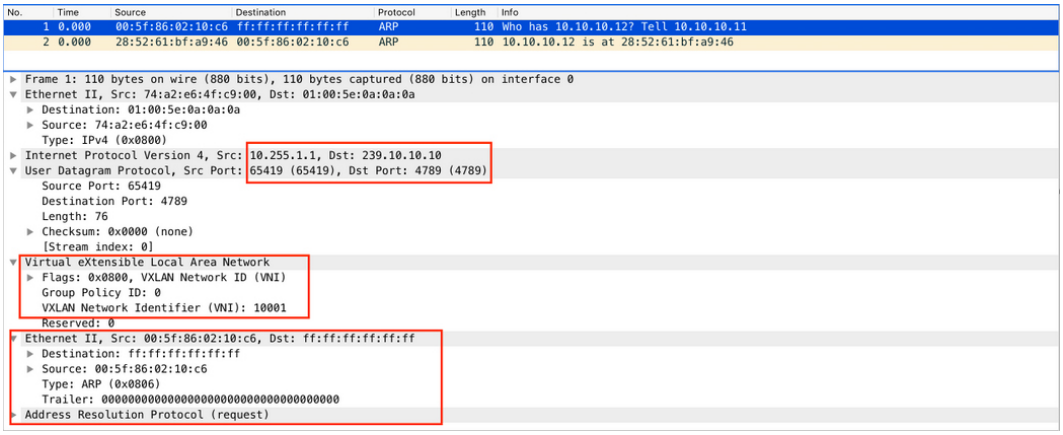
The following output confirms that the ARP request replies are present in embedded capture:

```
VTEP-1# show monitor capture 1 buffer display-filter "arp"
Starting the packet display ........ Press Ctrl + Shift + 6 to exit

7 0.000018 00:5f:86:02:10:c6 -> ff:ff:ff:ff:ff:ff ARP 110 Who has 10.10.10.12? Tell
10.10.10.11
9 0.000022 28:52:61:bf:a9:46 -> 00:5f:86:02:10:c6 ARP 110 10.10.10.12 is at 28:52:61:bf:a9:46
```

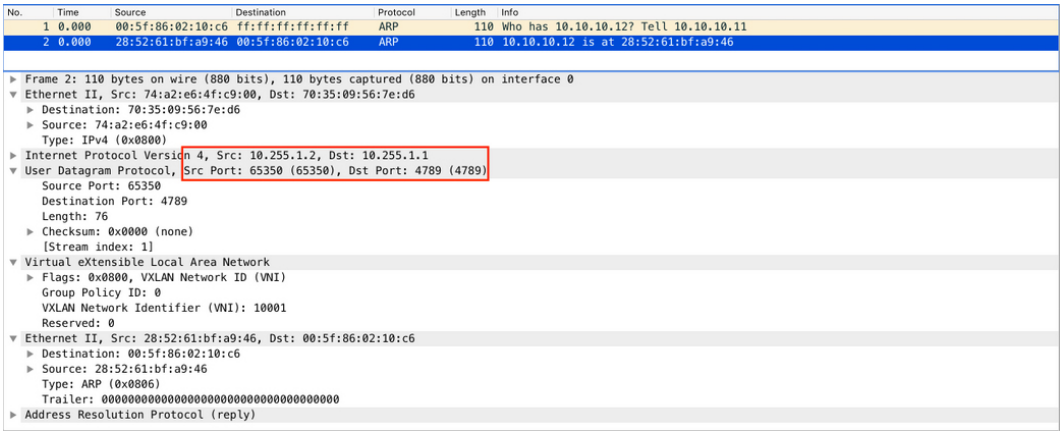
Verify that the Encapsulated ARP Request is Leaving in a Multicast Group to a VXLAN UDP Destination Port

The following image shows the ARP request leaving encapsulated in the multicast group 239.10.10.10, sourced from a VXLAN Loopback, to the VXLAN UDP destination port 4789 in the VNI 10001 and VLAN 10.



Verify that the ARP Reply from Core Interface is Encapsulated in Unicast to a VXLAN UDP Destination Port

The following image shows the ARP reply from core interface that is encapsulated in unicast, between VXLAN Loopbacks, to the VXLAN UDP destination port 4789 in the VNI 10001 and VLAN 10.



Once all of the above checks are verified, if there is still a problem with broadcast reachability, then repeat the checks on the remote VTEP.

Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI

This scenario might occur when host device 2 in VLAN 10 attempts to ping host device 3 that is also in VLAN 10. Perform the checks listed in the following table before troubleshooting unicast forwarding between VTEPs in the same VLAN through a Layer 2 VNI:

Table 58: Scenario 2: Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI

Check to be Performed	Steps to Follow
Has ARP been resolved on the local host for the Layer 2 adjacent remote host?	Run the **arp –a** command in privileged EXEC mode on the host device.
Do the hosts have the same subnet masks?	Perform any of the following steps: • Check the host device. • Check the SVI configuration on the VTEP.
Do you have the EVPN instance configured on your local VTEP?	Run the following commands in privileged EXEC mode on the VTEP: • **show run | section l2vpn** • **show run | section vlan config** • **show run interface nve** *interface-number*
Has the remote MAC address been learned in platform MATM in the same VLAN as the local host?	Run the **show platform software fed switch active matm macTable vlan** *vlan-id* command in privileged EXEC mode on the VTEP to check for the remote MAC addresses in the same VLAN.

To troubleshoot unicast forwarding between two VTEPs in the same VLAN using a Layer 2 VNI, follow these steps:

- Verify the provisioning of the EVPN VXLAN Layer 2 overlay network.
- Verify intra-subnet traffic movement in the EVPN VXLAN Layer 2 overlay network.

Verifying the Provisioning of an EVPN VXLAN Layer 2 Overlay Network

To verify the provisioning of an EVPN VXLAN Layer 2 overlay network, perform these checks:

1. Verify the Provisioning of the EVPN Instance in EVPN Manager, on page 657
2. Ensure that an NVE Peer is Present for the Layer 2 VNI, on page 659
3. Verify the Provisioning of the Layer 2 VNI in NVE Component, on page 659
4. Verify That the Layer 2 VNI VXLAN Tunnel Pseudoport is added to the Access VLAN in Layer 2 Forwarding Information Base (FIB), on page 660

Verify the Provisioning of the EVPN Instance in EVPN Manager

The following examples show how to verify that the EVPN instance is provisioned in the EVPN manager:

```
VTEP-1# show run | section l2vpn
l2vpn evpn instance 10 vlan-based
encapsulation vxlan
```

```
route-target export 10:1     <<— Import or export right route-targets

route-target import 10:2     <<— Import or export right route-targets


VTEP-1# show run | section vlan config
vlan configuration 10
member evpn-instance 10 vni 10001   <<— EVPN instance & VNI mapped to the VLAN

VTEP-1# show run interface nve1
interface nve1
source-interface Loopback999
host-reachability protocol bgp
member vni10001 mcast-group 239.10.10.10     <<— VNI added to NVE interface

VTEP-1# show run interface loopback 999
interface Loopback999
description VxLAN Loopback
ip address 10.255.1.1 255.255.255.255
```

Note Run the **show run** commands on VTEP 2 to verify its configuration, if required.

```
VTEP-1# show l2vpn evpn evi 10 detail <<— VLAN number and EVPN Instance number
                                          are not always the same, confirm which
                                          EVPN Instance maps to your VLAN
                                          with the show l2vpn evpn evi command
EVPN instance:     10 (VLAN Based)   <<— EVPN Instance number does map to the VLAN.
  RD:              10.1.1.1:10 (auto)
  Import-RTs:      10:2   <<— Importing VTEP-2 (if you are not seeing the prefix,
                              check configuration for the right import/export statement
                              under the l2vpn evpn instance)
  Export-RTs:      10:1
  Per-EVI Label:   none
  State:           Established
  Encapsulation:   vxlan
  Vlan:            10      <<— Layer 2 VLAN
    Ethernet-Tag:  0
    State:         Established <<— If State is not "Established", there
                                   could be a misconfiguration
    Core If:       Vlan99
    Access If:     Vlan10
    NVE If:        nve1
    RMAC:          7035.0956.7edd
    Core Vlan:     99
    L2 VNI:        10001           <<— Layer 2 VNI
    L3 VNI:        99999
    VTEP IP:       10.255.1.1
    MCAST IP:     239.10.10.10    <<— BUM Group for flooded traffic (Layer 2 learning, etc)

    VRF:           vxlan
    IPv4 IRB: Enabled
    IPv6 IRB: Enabled
    Pseudoports:
      GigabitEthernet1/0/1 service instance 10
   <<— Layer 2 Access pseudoport (combination of Layer 2 port and service instance)
```

Note If only a Layer 2 overlay network has been configured for bridging, then the `Core If` , `Access If`, `RMAC`, `Core BD`, `L3 VNI`, and `VRF` fields do not show any values as they are not set.

```
VTEP-2# show l2vpn evpn evi 10 detail
EVPN instance:     10 (VLAN Based)
RD:                10.2.2.2:10 (auto)
Import-RTs:        10:1       <<— Importing VTEP-1 route-target
Export-RTs:        10:2
Per-EVI Label:     none
State:             Established
Encapsulation:     vxlan
Vlan:              10         <<— Layer 2 VLAN
  Ethernet-Tag:    0
  State:           Established
  Core If:         Vlan99
  Access If:       Vlan10
  NVE If:          nve1
  RMAC:            7486.0bc4.b75d
  Core Vlan:       99
  L2 VNI:          10001      <<— Layer 2 VNI
  L3 VNI:          99999
  VTEP IP:         10.255.2.1
  MCAST IP:        239.10.10.10
  VRF:             vxlan
  IPv4 IRB: Enabled
  IPv6 IRB: Enabled
  Pseudoports:
   GigabitEthernet2/0/1 service instance 10
   <<— Layer 2 Access pseudoport (combination of Layer 2 port and service instance)
```

Ensure that an NVE Peer is Present for the Layer 2 VNI

The following examples show how to check if an NVE peer is present for the Layer 2 VNI:

```
VTEP-1# show nve peers vni 10001    <<— This VNI is learned from "show l2vpn evpn evi"
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10001    L2CP 10.255.2.1       2              10001      UP   N/A 00:01:03
   <<— Layer 2 Control Plane (L2CP) peer for the VNI is an indicator that this is
       Layer 2 forwarding
   <<— Interface NVE1, L2CP, egress VNI are shown, state is UP for a time of 00:01:03

VTEP-2# show nve peers vni 10001
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       10001    L2CP 10.255.1.1       3              10001      UP   N/A 00:47:2
   <<— Interface NVE1, L2CP, egress VNI are shown, state is UP for a time of 00:47:02
```

Verify the Provisioning of the Layer 2 VNI in NVE Component

The following example shows how to verify that the Layer 2 VNI is provisioned in the NVE component:

```
VTEP-1# show nve vni 10001 detail     <<— VNI 10001 is correlated to VLAN 10
                                          from show l2vpn evpn evi
Interface  VNI        Multicast-group VNI state  Mode  VLAN  cfg vrf
nve1       10001      239.10.10.10    Up         L2CP  10    CLI vxlan
```

```
   <<— state is UP, type is Layer 2 VNI (L2CP); VLAN 10 is mapped to VNI 10001

L2 VNI IPv6 IRB down reason:
BDI or associated L3 BDI's IPv6 addr un-configured
IPv6 topo_id disabled


L2CP VNI local VTEP info:      <<— Layer 2 VNI provisioning
VLAN: 10                       <<— Confirms that mapping is with VLAN 10
SVI if handler: 0x4D
Local VTEP IP: 10.255.1.1      <<— VxLAN Tunnel IP

Core IRB info:      <<— Layer 3 VPN provisioning (not required for troubleshooting
                        a scenario with pure Layer 2 VPN packet path
L3VNI: 99999
VRF name: vxlan
VLAN: 99
V4TopoID: 0x2
V6TopoID: 0xFFFF
Local VTEP IP: 10.255.1.1
SVI if handler: 0x50
SVI MAC: 7035.0956.7EDD

VNI Detailed statistics:
   Pkts In   Bytes In   Pkts Out  Bytes Out
         0          0 18158681548 27383291735556
```

Verify That the Layer 2 VNI VXLAN Tunnel Pseudoport is added to the Access VLAN in Layer 2 Forwarding Information Base (FIB)

The following examples show how to verify that the Layer 2 VXLAN tunnel pseudoport is added to the access VLAN in Layer 2 FIB:

```
VTEP-1# show l2fib bridge-domain 10 detail    <<— Bridge-domain will be same as VLAN number
Bridge Domain : 10
  Reference Count : 14
  Replication ports count : 2
  Unicast Address table size : 3
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 5109, Ports: 2

  VxLAN Information :
    VXLAN_DEC nv1:10001:239.10.10.10

  Port Information :
    BD_PORT   Gi1/0/1:10      <<— Pseudoport has been added to bridge-domain:
                                  (physical port + the BD number for the VLAN)
    VXLAN_REP nv1:10001:239.10.10.10        <<— VXLAN Replication group

  Unicast Address table information :
    008e.7391.1946  VXLAN_CP  L:10001:10.255.1.1 R:10001:10.255.2.1

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5109, Ports: 2
```

```
VTEP-2# show l2fib bridge-domain 10 detail
Bridge Domain : 10
  Reference Count : 15
  Replication ports count : 2
  Unicast Address table size : 4
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 5109, Ports: 2

  VxLAN Information :
    VXLAN_DEC nv1:10001:239.10.10.10

  Port Information :
    BD_PORT   Gi2/0/1:10      <<- Pseudoport has been added to bridge-domain:
                                  (physical port + the BD number for the VLAN)
    VXLAN_REP nv1:10001:239.10.10.10       <<- VXLAN replication group

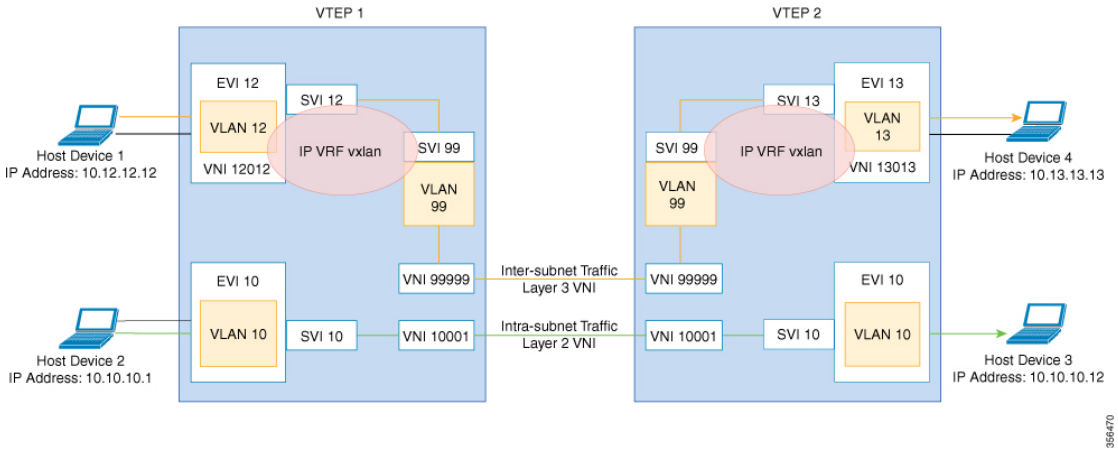
  Unicast Address table information :
    005f.8602.10c6  VXLAN_CP  L:10001:10.255.2.1 R:10001:10.255.1.1

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5109, Ports: 2
```

Verifying Intra-Subnet Traffic Movement in an EVPN VXLAN Layer 2 Overlay Network

The following figure illustrates the movement of traffic from host devices connected to VTEP 1 to host devices connected to VTEP 2:

Figure 61: Movement of traffic in an EVPN VXLAN network Through Layer 2 and Layer 3 VNIs



In the above figure, Layer 2 traffic moves from host device 2 to host device 3 through the Layer 2 VNI 10001. To verify the movement of intra-subnet traffic in the EVPN VXLAN Layer 2 overlay network, perform these checks:

1. Verify that the Local MAC Addresses Have Been Learned in IOS-MATM, on page 662
2. Verify that Both Local and Remote MAC Addresses are Learned in FED-MATM, on page 662

3. Confirm that the ICMP Echo Request Leaves VTEP 1 Encapsulated and Goes to a UDP Destination Port on VTEP 2, on page 663
4. Verify ARP for Local Host Devices, on page 663
5. Verify that the MAC Address Entries are Learned in SISF Device Tracking Table, on page 663
6. Verify that EVPN Manager Has Been Updated with the MAC Address Entries, on page 664
7. Verify that EVPN Manager Has Updated the MAC Routes into Layer 2 RIB, on page 665
8. Verify that Layer 2 RIB Has Updated BGP with the Local MAC Routes, and that BGP Has Updated Layer 2 RIB with the Remote MAC Routes, on page 665
9. Verify that the MAC Routes Learned from BGP and Updated to Layer 2 RIB are Also Updated to L2FIB, on page 667

Note Only MAC routes are considered while verifying the movement of intra-subnet traffic. MAC-IP routes are not applicable to bridged traffic.

Verify that the Local MAC Addresses Have Been Learned in IOS-MATM

The following examples show how to verify that the local MAC addresses have been learned in IOS-MATM:

```
VTEP-1# show mac address-table interface Gi1/0/1 vlan 10
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  10    005f.8602.10c6    DYNAMIC     Gi1/0/1    <<- IOS-MATM shows only
                                                      local MAC addresses

VTEP-2# show mac address-table interface g 2/0/1 vlan 10
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  10    008e.7391.1946    DYNAMIC     Gi2/0/1
```

Verify that Both Local and Remote MAC Addresses are Learned in FED-MATM

The following examples show how to verify that both local and remote MAC addresses are learned in FED-MATM:

```
VTEP-1# show platform software fed switch active matm macTable vlan 10
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle
siHandle           riHandle           diHandle             *a_time  *e_time  ports
-------------------------------------------------------------------------------------------------------------
10     005f.8602.10c6         0x1     60      0      0  0x7efcc0d78fc8      0x7efcc0ca8b88
      0x0                0x7efcc06cf9c8           300      144  GigabitEthernet1/0/1

   <<- Local MAC address is displayed here
```

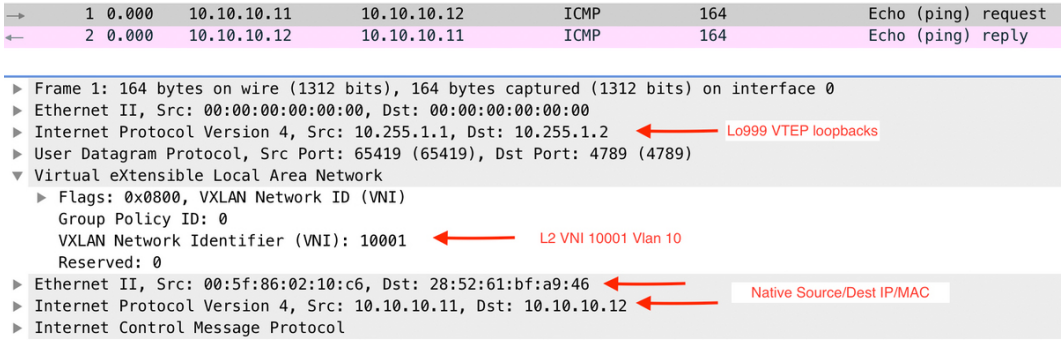
```
10     008e.7391.1946   0x1000001      0      0     64  0x7efcc0cafb38      0x7efcc0d7f628
      0x7ffa48c850b8      0x7efcc038cc18              0      144  RLOC 10.255.2.1 adj_id
135
   <<— Remote MAC address is displayed here

VTEP-2#sh platform software fed switch active matm macTable vlan 10
VLAN    MAC                  Type  Seq#   EC_Bi  Flags  machandle           siHandle
      riHandle            diHandle             *a_time  *e_time  ports
------------------------------------------------------------------------------------------
10     005f.8602.10c6   0x1000001      0      0     64  0x7fcec4e977d8      0x7fcec4e93ae8
      0x7fcec4e93308      0x7fcec430a3d8              0        0  RLOC 10.255.1.1 adj_id
64
   <<— Remote MAC address is displayed here
10     008e.7391.1946         0x1     46      0      0  0x7fcec4c6a248      0x7fcec4c20698
      0x0                 0x7fcec4611438            300      126  GigabitEthernet1/0/1

   <<— Local MAC address is displayed here
```

Confirm that the ICMP Echo Request Leaves VTEP 1 Encapsulated and Goes to a UDP Destination Port on VTEP 2

The following image confirms that the ICMP echo request leaves VTEP 1 encapsulated and goes to a UDP destination port on VTEP 2 through the loopback interface Lo999 and the Layer 2 VNI 10001:



Verify ARP for Local Host Devices

The following examples show how to verify ARP for local host devices:

```
VTEP-1# show ip arp vrf vxlan 10.10.10.11
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.10.10.11             2   005f.8602.10c6  ARPA   Vlan10

VTEP-2# show ip arp vrf vxlan 10.10.10.12
Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.10.10.12             4   008e.7391.1946  ARPA   Vlan10
```

Verify that the MAC Address Entries are Learned in SISF Device Tracking Table

The following examples show how to verify that the MAC addresses are learned in SISF device tracking table:

```
VTEP-1# show device-tracking database mac   <<— Only Local MAC addresses are seen
```

```
                                                  in SISF device tracking table
 MAC              Interface      vlan prlvl      state          time left policy
 005f.8602.10c6   Gi1/0/1          10 NO TRUST   MAC-REACHABLE  347 s     evpn-sisf-policy
   <<— MAC, REACH, and EVPN type SISF policy are displayed

VTEP-2# show device-tracking database mac   <<— Only Local MAC addresses are seen
                                                  in SISF device tracking table
 MAC              Interface      vlan prlvl      state          time left policy
 008e.7391.1946   Gi2/0/1          10 NO TRUST   MAC-REACHABLE  164 s     evpn-sisf-policy
   <<— MAC, REACH, and EVPN type SISF policy are displayed
```

Verify that EVPN Manager Has Been Updated with the MAC Address Entries

EVPN manager learns local MAC addresses and adds them to Layer 2 RIB. EVPN Manager also learns the remote MAC addresses from Layer 2 RIB, but the entries are only used for processing MAC mobility.

The following examples show how to verify that EVPN manager has been updated with the MAC addresses:

```
VTEP-1# show l2vpn evpn mac evi 10
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop
-------------- ----- ----- ---------------------------- ---------- ---------------
005f.8602.10c6 10    10    0000.0000.0000.0000.0000 0          Gi1/0/1
   <<— MAC Addresss learned by EVPN Manager. States look correct
008e.7391.1946 10    10    0000.0000.0000.0000.0000 0          10.255.2.1

VTEP-1#sh l2vpn evpn mac evi 10 detail
MAC Address:               005f.8602.10c6       <<— Local MAC address
EVPN Instance:             10      <<— EVPN Instance
Vlan:                      10      <<— VLAN
Ethernet Segment:          0000.0000.0000.0000.0000
Ethernet Tag ID:           0
Next Hop(s):               GigabitEthernet1/0/1 service instance 10<<— Local interface
                                                                      or local instance
VNI:                       10001      <<— VNI Label
Sequence Number:           0
MAC only present:          Yes
MAC Duplication Detection: Timer not running

MAC Address:               008e.7391.1946      <<— Remote MAC Address
EVPN Instance:             10       <<— EVPN Instance
Vlan:                      10       <<— VLAN
Ethernet Segment:          0000.0000.0000.0000.0000
Ethernet Tag ID:           0
Next Hop(s):               10.255.2.1      <<— Remote VTEP-2 Tunnel Loopback
Local Address:             10.255.1.1      <<— Local VTEP-1 Tunnel Loopback
VNI:                       10001       <<— VNI Label
Sequence Number:           0
MAC only present:          Yes
MAC Duplication Detection: Timer not running

VTEP-2# show l2vpn evpn mac evi 10
MAC Address    EVI   VLAN  ESI                          Ether Tag  Next Hop
-------------- ----- ----- ---------------------------- ---------- ---------------
005f.8602.10c6 10    10    0000.0000.0000.0000.0000 0          10.255.1.1
008e.7391.1946 10    10    0000.0000.0000.0000.0000 0          Gi2/0/1:10

VTEP-2#sh l2vpn evpn mac evi 10 detail
MAC Address:               005f.8602.10c6     <<— Remote MAC address
```

```
EVPN Instance:              10      <<— EVPN Instance
Vlan:                       10      <<— VLAN
Ethernet Segment:           0000.0000.0000.0000.0000
Ethernet Tag ID:            0
Next Hop(s):                10.255.1.1     <<— Remote VTEP-1 Tunnel Loopback
Local Address:              10.255.2.1     <<— Local VTEP-2 Tunnel Loopback
VNI:                        10001          <<— VNI Label
Sequence Number:            0
MAC only present:           Yes
MAC Duplication Detection: Timer not running

MAC Address:                008e.7391.1946       <<— Remote MAC adress
EVPN Instance:              10       <<— EVPN Instance
Vlan:                       10       <<— VLAN
Ethernet Segment:           0000.0000.0000.0000.0000
Ethernet Tag ID:            0
Next Hop(s):                GigabitEthernet2/0/1 service instance 10  <<— Local interface
                                                                          or local instance
VNI:                        10001      <<— VNI Label
Sequence Number:            0
MAC only present:           Yes
MAC Duplication Detection: Timer not running
```

Verify that EVPN Manager Has Updated the MAC Routes into Layer 2 RIB

Layer 2 RIB learns local MAC addresses from EVPN manager and updates BGP and Layer 2 FIB with them. Layer 2 RIB also learns remote MAC addresses from BGP and updates EVPN manager and Layer 2 FIB with them. Layer 2 RIB needs both local and remote MAC addresses in order to update BGP and Layer 2 FIB.

The following examples show how to verify that EVPN manager has updated the MAC routes into Layer 2 RIB:

```
VTEP-1# show l2route evpn mac
  EVI       ETag  Prod    Mac Address                Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------
   10          0 L2VPN 005f.8602.10c6                Gi1/0/1:10          0
   <<— Local prefix was added by EVPN Manager (Layer 2 VPN) into Layer 2 RIB
   10          0   BGP 008e.7391.1946          V:10001 10.255.2.1          0
   <<— Remote prefix was added by BGP into Layer 2 RIB

VTEP-2# show l2route evpn mac
  EVI       ETag  Prod    Mac Address                Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------
   10          0   BGP 005f.8602.10c6          V:10001 10.255.1.1          0
   <<— Remote prefix was added by BGP into Layer 2 RIB
   10          0 L2VPN 008e.7391.1946                Gi2/0/1:10          0
   <<— Local prefix was added by EVPN Manager (Layer 2 VPN) into Layer 2 RIB
```

Verify that Layer 2 RIB Has Updated BGP with the Local MAC Routes, and that BGP Has Updated Layer 2 RIB with the Remote MAC Routes

The following examples show how top verify that Layer 2 RIB has updated BGP with the local MAC routes and that BGP has updated Layer 2 RIB with the remote MAC routes:

```
VTEP-1# show bgp l2vpn evpn route-type 2 0 005f860210c6 *
   <<— Route-type is 2, Ethernet tag = 0, Local MAC address is in
       undelimited format, and * specifies to omit IP address
```

```
BGP routing table entry for [2][10.1.1.1:10][0][48][005F860210C6][0][*]/20, version 249
Paths: (1 available, best #1, table evi_10) <<— Added to BGP from EVPN Manager
                                                provisioning in l2vpn evi context
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.1.1.1)    <<— Locally Advertised by VTEP-1,
                                                     (:: indicates local)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 10001   <<— VNI ID is 10001 for VLAN 10
      Extended Community: RT:10:1 ENCAP:8    <<— RT 10:1 (local RT), Encap type 8 is VXLAN
      Local irb vxlan vtep:
        vrf:vxlan, l3-vni:99999
        local router mac:7035.0956.7EDD
        core-irb interface:Vlan99
        vtep-ip:10.255.1.1
      rx pathid: 0, tx pathid: 0x0


VTEP-1# show bgp l2vpn evpn route-type 2 0 008e73911946 *
   <<— Route-type is 2, Ethernet tag = 0, Remote MAC address is in
       undelimited format, and * specifies to omit IP address
BGP routing table entry for [2][10.1.1.1:10][0][48][008e73911946][0][*]/20, version 253
Paths: (1 available, best #1, table evi_10)    <<— EVPN instance BGP table for VLAN 10
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from [2][10.2.2.2:10][0][48][008e73911946][0][*]/20 (global)
   <<— From VTEP-2, RD is 10.2.2.2:10, MAC length is 48, [*] indicates MAC only
    10.255.2.1 (metric 2) (via default) from 10.2.2.2 (10.2.2.2)
   <<— Next hop of VTEP-2 Lo999, learned from RR 10.2.2.2
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10001    <<— VNI ID 10001 for VLAN 10
      Extended Community: RT:10:2 ENCAP:8     <<— Layer 2 VPN Route-Target 10:2
                                                   Encap type 8 is VXLAN
      Originator: 10.2.2.2, Cluster list: 10.2.2.2
      rx pathid: 0, tx pathid: 0x0

BGP routing table entry for [2][10.2.2.2:10][0][48][008e73911946][0][*]/20, version 251
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  Local
    10.255.2.1 (metric 2) (via default) from 10.2.2.2 (10.2.2.2)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10001
      Extended Community: RT:10:2 ENCAP:8
      Originator: 10.2.2.2, Cluster list: 10.2.2.2
      rx pathid: 0, tx pathid: 0x0


VTEP-2# show bgp l2vpn evpn route-type 2 0 008e73911946 *
   <<— Route-type is 2, Ethernet tag = 0, Local MAC address is in
       undelimited format, and * specifies to omit IP address
BGP routing table entry for [2][10.2.2.2:10][0][48][008e73911946][0][*]/20, version 292
Paths: (1 available, best #1, table evi_10)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (10.2.2.2)    <<— Locally Advertised by VTEP-2,
                                                     (:: indicates local)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 10001    <<— VNI ID 10001 for VLAN 10
```

```
      Extended Community: RT:10:2 ENCAP:8    <<— RT 10:2 (local RT), Encap type 8 is VXLAN
      Local irb vxlan vtep:
        vrf:vxlan, l3-vni:99999
        local router mac:7486.0BC4.B75D
        core-irb interface:Vlan99
        vtep-ip:10.255.2.1
      rx pathid: 0, tx pathid: 0x0


VTEP-2# show bgp l2vpn evpn route-type 2 0 005f860210c6 *
   <<— Route-type is 2, Ethernet tag = 0, Remote MAC address is in
       undelimited format, and * specifies to omit IP address
BGP routing table entry for [2][10.1.1.1:10][0][48][005F860210C6][0][*]/20, version 312
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 7
  Local
    10.255.1.1 (metric 2) (via default) from 10.2.2.2 (10.2.2.2)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10001
      Extended Community: RT:10:1 ENCAP:8
      Originator: 10.1.1.1, Cluster list: 10.2.2.2
      rx pathid: 0, tx pathid: 0x0

BGP routing table entry for [2][10.2.2.2:10][0][48][005F860210C6][0][*]/20, version 314
Paths: (1 available, best #1, table evi_10)    <<— EVPN instance BGP table for VLAN 10
  Not advertised to any peer
  Refresh Epoch 7
  Local, imported path from [2][10.1.1.1:10][0][48][005F860210C6][0][*]/20 (global)
   <<— From VTEP-2, RD is 10.2.2.2:10, MAC length is 48, [*] indicates MAC only
   <<— From VTEP-1, RD is 10.1.1.1:10, MAC length is 48, [*] indicates MAC only
    10.255.1.1 (metric 2) (via default) from 10.2.2.2 (10.2.2.2)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10001     <<— VNI ID 10001 for VLAN 10
      Extended Community: RT:10:1 ENCAP:8     <<— Layer 2 VPN Route-Target 10:1
                                                   Encap type 8 is VXLAN
      Originator: 10.1.1.1, Cluster list: 10.2.2.2
      rx pathid: 0, tx pathid: 0x0
```

Verify that the MAC Routes Learned from BGP and Updated to Layer 2 RIB are Also Updated to L2FIB

The following examples show how to verify that the MAC routes that are learned from BGP and updated to Layer 2 RIB are also updated to Layer 2 FIB:

```
VTEP-2# show l2fib bridge-domain 10 detail
Bridge Domain : 10
  Reference Count : 15
  Replication ports count : 2
  Unicast Address table size : 4
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 5109, Ports: 2

  VxLAN Information :
    VXLAN_DEC nv1:10001:239.10.10.10

  Port Information :
    BD_PORT   Gi2/0/1:10
    VXLAN_REP nv1:10001:239.10.10.10
```

```
  Unicast Address table information :
    005f.8602.10c6  VXLAN_CP  L:10001:10.255.2.1 R:10001:10.255.1.1
   <<— Remote MAC address is learned (local MAC address is not expected to be present)

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5109, Ports: 2

VTEP-1# show l2fib bridge-domain 10 detail
Bridge Domain : 10
  Reference Count : 14
  Replication ports count : 2
  Unicast Address table size : 3
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 5109, Ports: 2

  VxLAN Information :
    VXLAN_DEC nv1:10001:239.10.10.10

  Port Information :
    BD_PORT   Gi1/0/1
    VXLAN_REP nv1:10001:239.10.10.10

  Unicast Address table information :
    008e.7391.1946  VXLAN_CP  L:10001:10.255.1.1 R:10001:10.255.2.1
   <<— Remote MAC address is learned (local MAC address is not expected to be present)

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5109, Ports: 2
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5109, Ports: 2
```

Note Only remote MAC routes are displayed in the output.

Troubleshooting Unicast Forwarding Between VTEPS in Different VLANs Through a Layer 3 VNI

This scenario might occur when host device 1 in VLAN 12 attempts to ping host device 4 in VLAN 13. Perform the checks listed in the following table before troubleshooting unicast forwarding between VTEPs in different VLANs through a Layer 3 VNI:

Table 59: Scenario 3: Troubleshooting Unicast Forwarding Between VTEPS in Different VLANs Through a Layer 3 VNI

Check to be Performed	Steps to Follow
Are the source and destination host devices in different subnets?	Check the subnet of the local host device and compare it against the subnet of the remote host device.

Check to be Performed	Steps to Follow
Do you have an SVI interface configured for the remote subnet?	Run the **show ip interface brief | excluse unassigned** command in privileged EXEC mode on the VTEP.
Do you have the EVPN instance configured on your local VTEP?	Run the following commands in privileged EXEC mode on the VTEP: • **show run | section l2vpn** • **show run | section vlan config** • **show run interface nve** *interface-number*

To troubleshoot unicast forwarding between two VTEPs in different VLANs using a Layer 3 VNI, follow these steps:

- Verify the provisioning of the EVPN VXLAN Layer 3 overlay network.
- Verify inter-subnet traffic movement and symmetric IRB in the EVPN VXLAN Layer 3 overlay network.

Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network

To verify the provisioning of an EVPN VXLAN Layer 3 overlay network, perform these checks:

1. Verify that the Access SVIs, Core SVIs, and NVE Interfaces are Up, on page 669
2. Verify that the IP VRF is Provisioned with the Correct SVIs, Stitching Route-Targets, and Route Distinguisher, on page 670
3. Verify that Both Layer 2 and Layer 3 VNIs are provisioned in the VRF and are UP, on page 671
4. Verify that EVPN Manager is Updated from the NVE with all the Layer 2 and IRB Attributes, on page 672
5. Verify that the Remote Layer 3 VNI Details are Learned on Each VTEP, on page 673
6. Verify that the Layer 3 VNI Tunnel Pseudoport is Installed into Layer 2 FIB in the Core VLAN, on page 673

Verify that the Access SVIs, Core SVIs, and NVE Interfaces are Up

The following examples show how to verify that the access SVIs, core SVIs, and NVE interfaces are up:

```
VTEP-1# show ip interface brief
Interface              IP-Address      OK? Method Status         Protocol
Vlan10                 10.10.10.1      YES NVRAM  up             up
Vlan12                 10.12.12.1      YES NVRAM  up             up  <<— Access Interface
Vlan99                 10.255.1.1      YES unset  up             up  <<— Core Interface
   <<— If protocol status for the core interface is down, run the no autostate command
Loopback0              10.1.1.1        YES NVRAM  up             up
Loopback999            10.255.1.1      YES NVRAM  up             up
Tunnel0                10.255.1.1      YES unset  up             up
```

```
Tunnel1                10.1.1.5        YES unset  up             up
nve1                   unassigned      YES unset  up             up


VTEP-2# show ip interface brief
Interface              IP-Address      OK? Method Status         Protocol
Vlan10                 10.10.10.1      YES NVRAM  up             up
Vlan13                 10.13.13.1      YES NVRAM  up             up  <<— Access Interface
Vlan99                 10.255.2.1      YES unset  up             up  <<— Core Interface
   <<— If protocol status for the core interface is down, run the no autostate command
Loopback0              10.2.2.2        YES NVRAM  up             up
Loopback999            10.255.2.1      YES NVRAM  up             up
Tunnel0                10.255.2.1      YES unset  up             up
Tunnel1                10.1.1.10       YES unset  up             up
```

Verify that the IP VRF is Provisioned with the Correct SVIs, Stitching Route-Targets, and Route Distinguisher

The following examples show how to verify that the IP VRF is provisioned with the correct SVIs, stitching route-targets, and route distinguisher:

```
VTEP-1# show run vrf vxlan    <<— vxlan is the name of the VRF
vrf definition vxlan
rd 10.255.1.1:1
!
address-family ipv4
 route-target export 10.255.1.1:1 stitching     <<— Exporting local route-target
 route-target import 10.255.2.1:1 stitching     <<— Importing VTEP-2 route-target
```

```
VTEP-1# show ip vrf vxlan    <<— vxlan is the name of the VRF
  Name                             Default RD            Interfaces
  vxlan                            10.255.1.1:1          Vl10
                                                         Vl12
                                                         Vl99
```

```
VTEP-1# show ip vrf detail vxlan    <<— vxlan is the name of the VRF
VRF vxlan (VRF Id = 2); default RD 10.255.1.1:1; default VPNID <not set>
New CLI format, supports multiple address-families
Flags: 0x180C
Interfaces:
Vl10 Vl12 Vl99
Address family ipv4 unicast (Table ID = 0x2):   <<— Table 2 maps to VRF vxlan,
                                                    also found in BPG VPNv4 table
Flags: 0x0
No Export VPN route-target communities
No Import VPN route-target communities
Export VPN route-target stitching communities
   <<— VRF is using stitching route-targets. VTEPs must
       import each other's targets (same as Layer 3 VPN)
RT:10.255.1.1:1
Import VPN route-target stitching communities
RT:10.255.2.1:1
No import route-map
No global export route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix
```

```
VTEP-2# show ip vrf vxlan    <<— vxlan is the name of the VRF
  Name                             Default RD            Interfaces
```

```
  vxlan                            10.255.2.1:1          Vl110
                                                         Vl113
                                                         Vl199


VTEP-2# show ip vrf detail vxlan    <<— vxlan is the name of the VRF
VRF vxlan (VRF Id = 2); default RD 10.255.2.1:1; default VPNID <not set>
New CLI format, supports multiple address-families
Flags: 0x180C
Interfaces:
Vl10 Vl13 Vl199
Address family ipv4 unicast (Table ID = 0x2):   <<— Table 2 maps to VRF vxlan,
                                                     also found in BPG VPNv4 table
Flags: 0x0
No Export VPN route-target communities
No Import VPN route-target communities
Export VPN route-target stitching communities
   <<— VRF is using stitching route-targets. VTEPs must
       import each other's targets (same as Layer 3 VPN)
RT:10.255.2.1:1
Import VPN route-target stitching communities
RT:10.255.1.1:1
No import route-map
No global export route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix
```

Verify that Both Layer 2 and Layer 3 VNIs are provisioned in the VRF and are UP

The following examples show how to verify that both Layer 2 and Layer 3 VNIs are provisioned in the VRF and are up:

```
VTEP-1# show run | section vlan config
vlan configuration 99    <<— VNI is a member of VRF vxlan, not of EVPN instance
 member vni99999


VTEP-1# show run interface vlan 99
interface Vlan99
 description connected to L3_VNI_99999
 vrf forwarding vxlan
 ip unnumbered Loopback999


VTEP-1# show run interface nve 1
no ip address
 source-interface Loopback999
 host-reachability protocol bgp
 member vni 99999 vrf vxlan     <<— VNI tied to the VRF under NVE interface
 member vni 12012 mcast-group 239.12.12.12 <<— VNI tied to the NVE


VTEP-1# show run | section l2vpn
l2vpn evpn instance 12 vlan-based
 encapsulation vxlan
 route-target export 12:1    <<— Remote VTEP is NOT importing this route target,
                                 as it does not have the VLAN or VNI on its end
 route-target import 12:1
 no auto-route-target


VTEP-1# show run | section vlan config
```

```
vlan configuration 12
 member evpn-instance 12 vni 12012  <<— EVPN instance or VNI associated to the VLAN

VTEP-1# show nve vni
Interface  VNI        Multicast-group VNI state  Mode  VLAN  cfg vrf
nve1       10001      239.10.10.10    Up         L2CP  10    CLI vxlan
nve1       12012      239.12.12.12    Up         L2CP  12    CLI vxlan <<— Layer 2 VNI
nve1       99999      N/A             Up         L3CP  99    CLI vxlan <<— Layer 3 VNI

VTEP-2# show nve vni
Interface  VNI        Multicast-group VNI state  Mode  VLAN  cfg vrf
nve1       13013      239.13.13.13    Up         L2CP  13    CLI vxlan <<— Layer 2 VNI
nve1       10001      239.10.10.10    Up         L2CP  10    CLI vxlan
nve1       99999      N/A             Up         L3CP  99    CLI vxlan <<— Layer 3 VNI
```

Verify that EVPN Manager is Updated from the NVE with all the Layer 2 and IRB Attributes

The following examples show how to verify that EVPN manager is updated from the NVE with all the Layer 2 and IRB attributes:

```
VTEP-1# show l2vpn evpn evi
EVI   VLAN  Ether Tag  L2 VNI    Multicast     Pseudoport
----- ----- ---------- --------- ------------- ------------------
12    12    0          12012     239.12.12.12  Gi1/0/1
   <<— See which EVPN instance maps to the VLAN. The VLAN
       or EVPN instance values are not always the same
<...snip...>

VTEP-1# show l2vpn evpn evi 12 detail
EVPN instance:    12 (VLAN Based)
  RD:             10.1.1.1:12 (auto)
  Import-RTs:     12:1
  Export-RTs:     12:1
  Per-EVI Label:  none
  State:          Established
  Encapsulation:  vxlan
  Vlan:           12    <<— VLAN Layer 2 VNI
    Ethernet-Tag: 0
    State:        Established
    Core If:      Vlan99    <<— Interface handling IP VRF forwarding
    Access If:    Vlan12
    NVE If:       nve1
    RMAC:         7035.0956.7edd  <<— RMAC is the BIA of SVI 99 Core interface
    Core Vlan:    99
    L2 VNI:       12012
    L3 VNI:       99999
    VTEP IP:      10.255.1.1   <<— Local Tunnel endpoint IP address
    MCAST IP:     239.12.12.12
    VRF:          vxlan    <<— IP VRF for Layer 3 VPN
    Pseudoports:
      GigabitEthernet1/0/1 service instance 12

VTEP-2# show l2vpn evpn evi
EVI   VLAN  Ether Tag  L2 VNI    Multicast     Pseudoport
----- ----- ---------- --------- ------------- ------------------
13    13    0          13013     239.13.13.13  Gi2/0/1:13
   <<— See which EVPN instance maps to the VLAN. The VLAN
       or EVPN instance values are not always the same
```

```
VTEP-2# show l2vpn evpn evi 13 detail
EVPN instance:    13 (VLAN Based)
  RD:             10.2.2.2:13 (auto)
  Import-RTs:     13:2
  Export-RTs:     13:2
  Per-EVI Label:  none
  State:          Established
  Encapsulation:  vxlan
  Vlan:           13     <<— VLAN Layer 2 VNI
    Ethernet-Tag: 0
    State:        Established
    Core If:      Vlan99    <<— Interface handling IP VRF forwarding
    Access If:    Vlan13
    NVE If:       nve1
    RMAC:         7486.0bc4.b75d   <<— RMAC is the BIA of SVI 99 Core interface
    Core Vlan:    99
    L2 VNI:       13013
    L3 VNI:       99999
    VTEP IP:      10.255.2.1   <<— Local Tunnel endpoint IP address
    MCAST IP:     239.13.13.13
    VRF:          vxlan    <<— IP VRF for Layer 3 VPN
    Pseudoports:
      GigabitEthernet2/0/1 service instance 13
```

Verify that the Remote Layer 3 VNI Details are Learned on Each VTEP

The following examples show how to verify that the remote Layer 3 VNI details are learned on each VTEP:

```
VTEP-1# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       99999    L3CP 10.255.2.1       7486.0bc4.b75d 99999      UP   A/M 1w1d
   <<— Layer 3 Control Plane (L3CP), RMAC of Remote VTEP and Uptime of peer are displayed
```

```
VTEP-2# show nve peers
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags UP time
nve1       99999    L3CP 10.255.1.1       7035.0956.7edd 99999      UP   A/M 21:27:36
   <<— Layer 3 Control Plane (L3CP), RMAC of Remote VTEP and Uptime of peer are displayed
```

Verify that the Layer 3 VNI Tunnel Pseudoport is Installed into Layer 2 FIB in the Core VLAN

The following examples show how to verify that the Layer 3 VNI tunnel pseudoport is installed into Layer 2 FIB in the core VLAN:

```
VTEP-1# show l2fib bridge-domain 99 detail
   <<— The Core VLAN can be obtained in the output of the
       show l2vpn evpn evi <evpn-instance> detail command
Bridge Domain : 99
  Reference Count : 8
  Replication ports count : 0
  Unicast Address table size : 1
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 5112, Ports: 0
```

```
  VxLAN Information :

  Unicast Address table information :
    7486.0bc4.b75d  VXLAN_CP  L:99999:10.255.1.1 R:99999:10.255.2.1
   <<— Encapsulation Information to reach remote VTEP-2

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5112, Ports: 0
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5112, Ports: 0
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5112, Ports: 0


VTEP-2# show l2fib bridge-domain 99 detail
   <<— The Core VLAN can be obtained in the output of the
       show l2vpn evpn evi <evpn-instance> detail command

Bridge Domain : 99
  Reference Count : 8
  Replication ports count : 0
  Unicast Address table size : 1
  IP Multicast Prefix table size : 3

  Flood List Information :
    Olist: 5111, Ports: 0

  VxLAN Information :

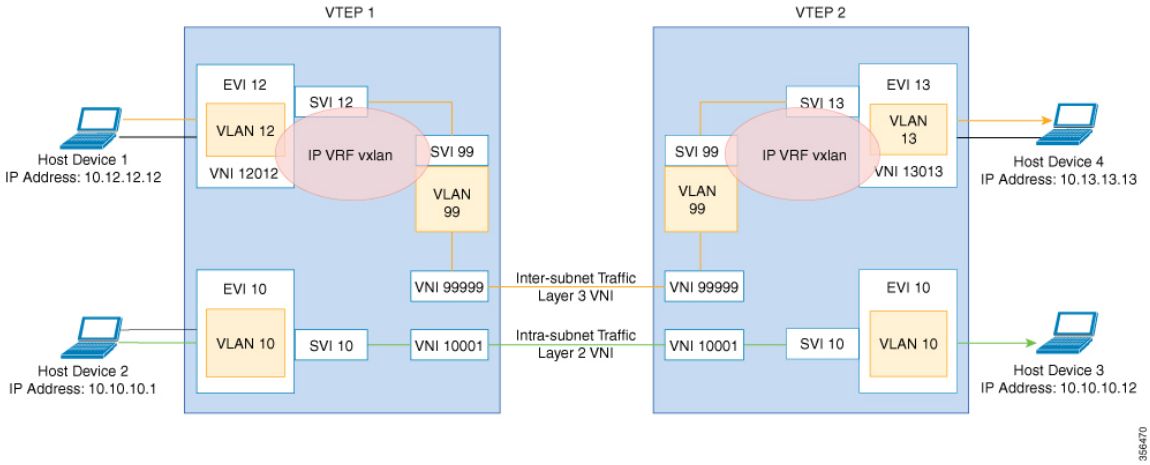
  Unicast Address table information :
    7035.0956.7edd  VXLAN_CP  L:99999:10.255.2.1 R:99999:10.255.1.1
   <<— Encapsulation Information to reach remote VTEP-2

  IP Multicast Prefix table information :
    Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5111, Ports: 0
    Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5111, Ports: 0
    Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5111, Ports: 0
```

Verifying Inter-Subnet Traffic Movement and Symmetric IRB in an EVPN VXLAN Layer 3 Overlay Network

The following figure illustrates the movement of traffic from host devices connected to VTEP 1 to host devices connected to VTEP 2:

Figure 62: Movement of traffic in an EVPN VXLAN network through Layer2 and Layer 3 VNIs



In the above figure, Layer 3 traffic moves from host device 1 to host device 4 through the Layer 3 VNI 99999. To verify the movement of inter-subnet traffic in the EVPN VXLAN Layer 3 overlay network, perform these checks:

1. Verify that Local MAC Address and IP Address Entries are Learned in SISF Device Tracking Table, on page 675
2. Verify that MAC Address and IP Address Entries are Learned in EVPN Manager, on page 676
3. Verify that MAC Address and IP Address Entries are Learned in Layer 2 RIB, on page 677
4. Verify that Local MAC Address and IP Address Entries are Learned in MAC VRF, on page 677
5. Verify that Remote MAC-IP Address Pair is Learend in the VRF, on page 678
6. Verify that IP Routes are Inserted in RIB, on page 679
7. Verify that the Adjacency Table Contains Entries for the VRF-Enabled Core VLAN Interface, on page 679
8. Confirm that Adjacency Exists to the VTEP Tunnel IP Address for a Host Device in IP VRF, on page 680
9. Confirm that Adjacency Exists to Reach Tunnel Destination, on page 680
10. Confirm that the ICMP Echo Request that Leaves Encapsulated from the Source VTEP Reaches the Loopback Tunnel Endpoint and UDP Destination Port on the Destination VTEP Through the Layer 3 VNI and IP VRF, on page 680

Verify that Local MAC Address and IP Address Entries are Learned in SISF Device Tracking Table

The following examples show how to verify that local MAC address and IP address entries are learned in SISF device tracking table:

```
VTEP-1# show device-tracking database vlanid 12
Binding Table has 4 entries, 2 dynamic (limit 100000)
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol,
DH4 - IPv4 DHCP, DH6 - IPv6 DHCP, PKT - Other Packet, API - API created
```

```
Preflevel flags (prlvl):
0001:MAC and LLA match     0002:Orig trunk            0004:Orig access
0008:Orig trusted trunk    0010:Orig trusted access   0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated    0100:Statically assigned


    Network Layer Address              Link Layer Address Interface        vlan prlvl  age
   state     Time left
ARP 10.12.12.12                            005f.8602.10e7  Gi1/0/1          12  0005  115s
  REACHABLE  N/A
```

```
VTEP-2# show device-tracking database vlanid 13
vlanDB has 2 entries for vlan 13, 1 dynamic
Codes: L - Local, S - Static, ND - Neighbor Discovery, ARP - Address Resolution Protocol,
DH4 - IPv4 DHCP, DH6 - IPv6 DHCP, PKT - Other Packet, API - API created
Preflevel flags (prlvl):
0001:MAC and LLA match     0002:Orig trunk            0004:Orig access
0008:Orig trusted trunk    0010:Orig trusted access   0020:DHCP assigned
0040:Cga authenticated     0080:Cert authenticated    0100:Statically assigned


    Network Layer Address              Link Layer Address Interface        vlan prlvl  age
   state     Time left
ARP 10.13.13.13                            008e.7391.1977  Gi2/0/1          13  0005  155s
  REACHABLE  N/A
```

Verify that MAC Address and IP Address Entries are Learned in EVPN Manager

The following examples show how to verify that MAC address and IP address entries are learned in EVPN manager:

```
VTEP-1# show l2vpn evpn mac ip evi 12
IP Address                EVI   VLAN  MAC Address    Next Hop
------------------------ ----- ----- -------------- -------------------------
10.12.12.12               12    12    005f.8602.10e7 Gi1/0/1

VTEP-1#sh l2vpn evpn mac ip evi 12 detail
IP Address:                10.12.12.12
EVPN Instance:             12
Vlan:                      12
MAC Address:               005f.8602.10e7
Ethernet Segment:          0000.0000.0000.0000.0000
Ethernet Tag ID:           0
Next Hop:                  GigabitEthernet1/0/1 service instance 12
VNI:                       12012
Sequence Number:           0
IP Duplication Detection:  Timer not running


VTEP-2# show l2vpn evpn mac ip evi 13
IP Address                EVI   VLAN  MAC Address    Next Hop
------------------------ ----- ----- -------------- -------------------------
10.13.13.13               13    13    008e.7391.1977 Gi2/0/1:13

VTEP-2#sh l2vpn evpn mac ip evi 13 detail
IP Address:                10.13.13.13
EVPN Instance:             13
Vlan:                      13
MAC Address:               008e.7391.1977
Ethernet Segment:          0000.0000.0000.0000.0000
```

```
Ethernet Tag ID:           0
Next Hop:                  GigabitEthernet2/0/1 service instance 13
VNI:                       13013
Sequence Number:           0
IP Duplication Detection:  Timer not running
```

Verify that MAC Address and IP Address Entries are Learned in Layer 2 RIB

The following examples show how to verify that MAC address and IP address entries are learned in Layer 2 RIB:

```
VTEP-1# show l2route evpn mac ip
  EVI       ETag  Prod    Mac Address         Host IP                Next Hop(s)
----- ---------- ----- -------------- --------------- -------------------------
   12          0 L2VPN 005f.8602.10e7     10.12.12.12                Gi1/0/1

VTEP-2# show l2route evpn mac ip
  EVI       ETag  Prod    Mac Address         Host IP                Next Hop(s)
----- ---------- ----- -------------- --------------- -------------------------
   13          0 L2VPN 008e.7391.1977     10.13.13.13                Gi1/0/1
```

Verify that Local MAC Address and IP Address Entries are Learned in MAC VRF

```
VTEP-1# show bgp l2vpn evpn evi 12 route-type 2 0 005F860210E7 10.12.12.12
BGP routing table entry for [2][10.1.1.1:12][0][48][005F860210E7][32][10.12.12.12]/24,
version 72
Paths: (1 available, best #1, table evi_12)    <<— The Layer 2 VPN table number
                                                    for EVPN instance 12
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local    <<— Indicates locally learned route
    :: (via default) from 0.0.0.0 (10.1.1.1)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 12012, Label2 99999 <<— Displays both Layer 2
                                                                     and VRF labels
      Extended Community: RT:12:1 RT:10.255.1.1:1 ENCAP:8   <<— Note the VRF stitching RT
                                                                as well as the Layer 2 RT
        Router MAC:7035.0956.7EDD
      Local irb vxlan vtep:
        vrf:vxlan, l3-vni:99999
        local router mac:7035.0956.7EDD   <<— Local RMAC
        core-irb interface:Vlan99    <<— VRF Layer 3 VPN interface
        vtep-ip:10.255.1.1   <<— Loopback 999 tunnel endpoint
      rx pathid: 0, tx pathid: 0x0
```

The following examples show how to verify that local MAC address and IP address entries are learned in MAC VRF:

```
VTEP-2# show bgp l2vpn evpn evi 13 route-type 2 0 008E73911977 10.13.13.13
BGP routing table entry for [2][10.2.2.2:13][0][48][008E73911977][32][10.13.13.13]/24,
version 70
Paths: (1 available, best #1, table evi_13)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local    <<— Indicates locally learned route
```

```
    :: (via default) from 0.0.0.0 (10.2.2.2)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
      EVPN ESI: 00000000000000000000, Label1 13013, Label2 99999
      Extended Community: RT:13:1 RT:10.255.2.1:1 ENCAP:8
        Router MAC:7486.0BC4.B75D
      Local irb vxlan vtep:
        vrf:vxlan, l3-vni:99999
        local router mac:7486.0BC4.B75D
        core-irb interface:Vlan99
        vtep-ip:10.255.2.1
      rx pathid: 0, tx pathid: 0x0
```

Verify that Remote MAC-IP Address Pair is Learend in the VRF

The following examples verify that remote MAC-IP address pair is learned in the VRF:

```
VTEP-1# show bgp vpnv4 unicast vrf vxlan 10.13.13.13
BGP routing table entry for 10.255.1.1:1:10.13.13.13/32, version 15
Paths: (1 available, best #1, table vxlan)         <<— VPNv4 VRF BGP table
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [2][10.2.2.2:13][0][48][008E73911977][32][10.13.13.13]/24
(global)
   <<— EVPN type-2, l2vpn RD 10.2.2.2:13, MAC and IP addresses
    10.255.2.1 (metric 3) (via default) from 10.2.2.2 (10.2.2.2)
   <<— Next hop 10.255.2.1, learned from RR 10.2.2.2
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 Router MAC:7486.0BC4.B75D
      Originator: 10.2.2.2, Cluster list: 10.2.2.2
      Local vxlan vtep:
        vrf:vxlan, vni:99999
        local router mac:7035.0956.7EDD
        encap:8
        vtep-ip:10.255.1.1
        bdi:Vlan99
      Remote VxLAN:
        Topoid 0x2(vrf vxlan)   <<— VRF vxlan (mapped to ID 2)
        Remote Router MAC:7486.0BC4.B75D  <<— VTEP-2 RMAC
        Encap 8    <<— VXLAN encap (type 8)
        Egress VNI 99999    <<— VRF VNI
        RTEP 10.255.2.1    <<— VTEP-2 Remote Tunnel Endpoint
      rx pathid: 0, tx pathid: 0x0


VTEP-2# show bgp vpnv4 unicast vrf vxlan 10.12.12.12
BGP routing table entry for 10.255.2.1:1:10.12.12.12/32, version 15
Paths: (1 available, best #1, table vxlan)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [2][10.1.1.1:12][0][48][005F860210E7][32][10.12.12.12]/24
(global)
   <<— EVPN type-2, l2vpn RD 10.1.1.1:12, MAC and IP addresses
    10.255.1.1 (metric 3) (via default) from 10.2.2.2 (10.2.2.2)
   <<— Next hop 10.255.1.1, learned from RR 10.2.2.2
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 Router MAC:7035.0956.7EDD
      Originator: 10.1.1.1, Cluster list: 10.2.2.2
      Local vxlan vtep:
        vrf:vxlan, vni:99999
        local router mac:7486.0BC4.B75D
        encap:8
        vtep-ip:10.255.2.1
```

```
    bdi:Vlan99
  Remote VxLAN:
    Topoid 0x2(vrf vxlan)    <<— VRF vxlan (mapped to ID 2)
    Remote Router MAC:7035.0956.7EDD  <<— VTEP-1 RMAC
    Encap 8    <<— VXLAN encap (type 8)
    Egress VNI 99999    <<— VRF VNI
    RTEP 10.255.1.1    <<— VTEP-2 Remote Tunnel Endpoint
  rx pathid: 0, tx pathid: 0x0
```

Verify that IP Routes are Inserted in RIB

The following examples show how to verify that IP routes are inserted in RIB:

```
VTEP-1# show ip route vrf vxlan 10.13.13.13

Routing Table: vxlan
Routing entry for 10.13.13.13/32
  Known via "bgp 69420", distance 200, metric 0, type internal
  Last update from 10.255.2.1 on Vlan99, 00:11:33 ago
  Routing Descriptor Blocks:
  * 10.255.2.1 (default), from 10.2.2.2, 00:11:33 ago, via Vlan99 <<— Next hop here is the
                                                                    Core VLAN interface
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none


VTEP-2# show ip route vrf vxlan 10.12.12.12

Routing Table: vxlan
Routing entry for 10.12.12.12/32
  Known via "bgp 69420", distance 200, metric 0, type internal
  Last update from 10.255.1.1 on Vlan99, 00:04:06 ago
  Routing Descriptor Blocks:
  * 10.255.1.1 (default), from 10.2.2.2, 00:04:06 ago, via Vlan99 <<— Next hop here is the
                                                                    Core VLAN interface
      Route metric is 0, traffic share count is 1
      AS Hops 0
      MPLS label: none
```

Verify that the Adjacency Table Contains Entries for the VRF-Enabled Core VLAN Interface

The following examples show how to verify that the adjacency table contains entries for the VRF-enabled core VLAN interface:

```
VTEP-1# show adjacency vlan 99 detail
Protocol Interface                 Address
IP       Vlan99                    10.255.2.1(9)    <<— IP unnumbered from Loopback 999
                                   0 packets, 0 bytes
                                   epoch 0
                                   sourced in sev-epoch 6
                                   Encap length 14
                                   74860BC4B75D703509567EDD0800
   <<— Local RMAC is 74860BC4B75D, Remote RMAC is 703509567EDD, etype is 800
                                   VXLAN Transport tunnel
   <<— Tunnel Interface (RMAC, using VTEP Loopback IP address)

VTEP-2# show adjacency vlan 99 detail
```

```
Protocol Interface                 Address
IP       Vlan99                    10.255.1.1(9)   <<— IP unnumbered from Loopback 999
                                   0 packets, 0 bytes
                                   epoch 0
                                   sourced in sev-epoch 5
                                   Encap length 14
                                   703509567EDD74860BC4B75D0800
   <<— Local RMAC is 703509567EDD, Remote RMAC is 74860BC4B75D, etype is 800
                                   VXLAN Transport tunnel
   <<— Tunnel Interface (RMAC, using VTEP Loopback IP address)
```

Confirm that Adjacency Exists to the VTEP Tunnel IP Address for a Host Device in IP VRF

The following example shows how to confirm that adjacency exists to the VTEP Tunnel IP address for a host device in IP VRF:

```
VTEP-1# show ip cef vrf vxlan 10.13.13.13/32  <<— Remote host in VLAN 13 of VTEP-2
10.13.13.13/32
  nexthop 10.255.2.1 Vlan99
```

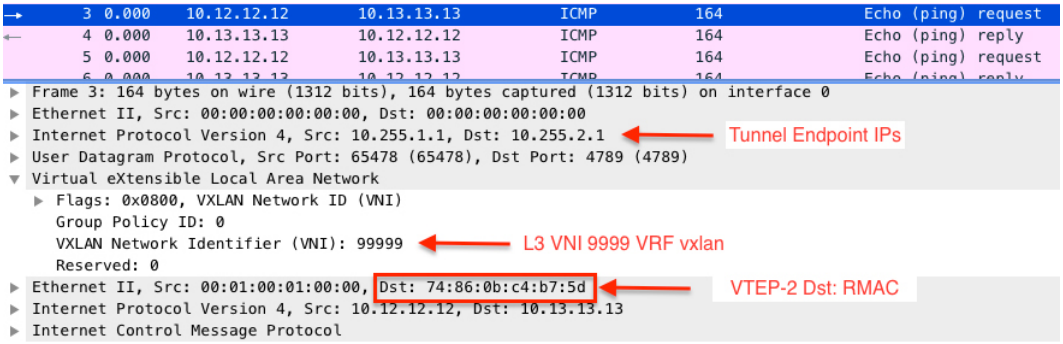
Confirm that Adjacency Exists to Reach Tunnel Destination

The following example shows how to confirm that adjacency exists to reach tunnel destination:

```
VTEP-1# show ip cef 10.255.1.11
10.255.2.1/32
  nexthop 10.1.1.6 GigabitEthernet1/0/1
```

Confirm that the ICMP Echo Request that Leaves Encapsulated from the Source VTEP Reaches the Loopback Tunnel Endpoint and UDP Destination Port on the Destination VTEP Through the Layer 3 VNI and IP VRF

The following image confirms that the ICMP echo request that leaves encapsulated from source VTEP reaches the Loopback interface and UDP destination port on the destination VTEP through the Layer 3 VNI and IP VRF:



Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network

This scenario might occur when host device 1 attempts to ping an external IP address through a border leaf VTEP. Perform the checks listed in the following table before troubleshooting unicast forwarding between a VXLAN network and an external IP network.

Table 60: Scenario 4: Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network

Check to be performed	Steps to follow
Is one IP address present in the VXLAN network and the other IP address coming from external IP network?	Check the local subnets (or the SVI interfaces) if the remote subnet is present. **Note** Local subnet has the remote subnet listed even in the case of scenario 3.
Is the EVPN route type 5 being used to send traffic to remote destination?	Run the **show bgp l2vpn evpn all** command in privileged EXEC mode on the VTEP. Look for remote prefix to be displayed as `[5]` for route type 5.

To troubleshoot unicast forwarding between a VXLAN network and an external IP network, follow these steps:

- Verify the provisioning of the EVPN VXLAN Layer 3 overlay network.
- Verify traffic movement from the VXLAN network to the IP network through the border leaf switch using route type 5.

Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network

Verify that the Access SVIs, Core SVIs, and NVE Interfaces are Up

The following examples show how to verify that the access SVIs, core SVIs, and NVE interfaces are up:

```
VTEP-1# show ip interface brief
Interface IP-Address OK? Method Status Protocol
Vlan10 10.10.10.1 YES NVRAM up up
Vlan12 10.12.12.1 YES NVRAM up up <<— Access Interface
Vlan99 10.255.1.1 YES unset up up <<— Core Interface
<<— If protocol status for the core interface is down, run the no autostate command
Loopback0 10.1.1.1 YES NVRAM up up
Loopback999 10.255.1.1 YES NVRAM up up
Tunnel0 10.255.1.1 YES unset up up
Tunnel1 10.1.1.5 YES unset up up
nve1 unassigned YES unset up up
VTEP-2# show ip interface brief
Interface IP-Address OK? Method Status Protocol
Vlan10 10.10.10.1 YES NVRAM up up
Vlan13 10.13.13.1 YES NVRAM up up <<— Access Interface
Vlan99 10.255.2.1 YES unset up up <<— Core Interface
<<— If protocol status for the core interface is down, run the no autostate command
```

```
Loopback0 10.2.2.2 YES NVRAM up up
Loopback999 10.255.2.1 YES NVRAM up up
Tunnel0 10.255.2.1 YES unset up up
Tunnel1 10.1.1.10 YES unset up up
```

Verify that the IP VRF is Provisioned with the Correct SVIs, Stitching Route-Targets, and Route Distinguisher

The following examples show how to verify that the IP VRF is provisioned with the correct SVIs, stitching route-targets, and route distinguisher:

```
VTEP-1# show run vrf vxlan <<— vxlan is the name of the VRF
vrf definition vxlan
rd 10.255.1.1:1
!
address-family ipv4
route-target export 10.255.1.1:1 stitching <<— Exporting local route-target
route-target import 10.255.2.1:1 stitching <<— Importing VTEP-2 route-target
VTEP-1# show ip vrf vxlan <<— vxlan is the name of the VRF
Name Default RD Interfaces
vxlan 10.255.1.1:1 Vl10
Vl12
Vl99
VTEP-1# show ip vrf detail vxlan <<— vxlan is the name of the VRF
VRF vxlan (VRF Id = 2); default RD 10.255.1.1:1; default VPNID <not set>
New CLI format, supports multiple address-families
Flags: 0x180C
Interfaces:
Vl10 Vl12 Vl99
Address family ipv4 unicast (Table ID = 0x2): <<— Table 2 maps to VRF vxlan,
also found in BPG VPNv4 table
Flags: 0x0
No Export VPN route-target communities
No Import VPN route-target communities
Export VPN route-target stitching communities
<<— VRF is using stitching route-targets. VTEPs must
import each other's targets (same as Layer 3 VPN)
RT:10.255.1.1:1
Import VPN route-target stitching communities
RT:10.255.2.1:1
No import route-map
No global export route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix
VTEP-2# show ip vrf vxlan <<— vxlan is the name of the VRF
Name Default RD Interfaces
vxlan 10.255.2.1:1 Vl10
Vl13
Vl99
VTEP-2# show ip vrf detail vxlan <<— vxlan is the name of the VRF
VRF vxlan (VRF Id = 2); default RD 10.255.2.1:1; default VPNID <not set>
New CLI format, supports multiple address-families
Flags: 0x180C
Interfaces:
Vl10 Vl13 Vl99
Address family ipv4 unicast (Table ID = 0x2): <<— Table 2 maps to VRF vxlan,
also found in BPG VPNv4 table
Flags: 0x0
No Export VPN route-target communities
No Import VPN route-target communities
Export VPN route-target stitching communities
<<— VRF is using stitching route-targets. VTEPs must
import each other's targets (same as Layer 3 VPN)
RT:10.255.2.1:1
```

```
Import VPN route-target stitching communities
RT:10.255.1.1:1
No import route-map
No global export route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix
```

Verify that Both Layer 2 and Layer 3 VNIs are provisioned in the VRF and are UP

The following examples show how to verify that both Layer 2 and Layer 3 VNIs are provisioned in the VRF and are up:

```
VTEP-1# show run | section vlan config
vlan configuration 99 <<— VNI is a member of VRF vxlan, not of EVPN instance
member vni99999
VTEP-1# show run interface vlan 99
interface Vlan99
description connected to L3_VNI_99999
vrf forwarding vxlan
ip unnumbered Loopback999
VTEP-1# show run interface nve 1
no ip address
source-interface Loopback999
host-reachability protocol bgp
member vni 99999 vrf vxlan <<— VNI tied to the VRF under NVE interface
member vni 12012 mcast-group 239.12.12.12 <<— VNI tied to the NVE
VTEP-1# show run | section l2vpn
l2vpn evpn instance 12 vlan-based
encapsulation vxlan
route-target export 12:1 <<— Remote VTEP is NOT importing this route target,
as it does not have the VLAN or VNI on its end
route-target import 12:1
no auto-route-target
VTEP-1# show run | section vlan config
vlan configuration 12
member evpn-instance 12 vni 12012 <<— EVPN instance or VNI associated to the VLAN
VTEP-1# show nve vni
Interface VNI Multicast-group VNI state Mode VLAN cfg vrf
nve1 10001 239.10.10.10 Up L2CP 10 CLI vxlan
nve1 12012 239.12.12.12 Up L2CP 12 CLI vxlan <<— Layer 2 VNI
nve1 99999 N/A Up L3CP 99 CLI vxlan <<— Layer 3 VNI
VTEP-2# show nve vni
Interface VNI Multicast-group VNI state Mode VLAN cfg vrf
nve1 13013 239.13.13.13 Up L2CP 13 CLI vxlan <<— Layer 2 VNI
nve1 10001 239.10.10.10 Up L2CP 10 CLI vxlan
nve1 99999 N/A Up L3CP 99 CLI vxlan <<— Layer 3 VNI
```

Verify that EVPN Manager is Updated from the NVE with all the Layer 2 and IRB Attributes

The following examples show how to verify that EVPN manager is updated from the NVE with all the Layer 2 and IRB attributes:

```
VTEP-1# show l2vpn evpn evi
EVI VLAN Ether Tag L2 VNI Multicast Pseudoport
----- ----- ---------- --------- ------------- ------------------
12 12 0 12012 239.12.12.12 Ge1/0/1
<<— See which EVPN instance maps to the VLAN. The VLAN
or EVPN instance values are not always the same
<...snip...>
VTEP-1# show l2vpn evpn evi 12 detail
EVPN instance: 12 (VLAN Based)
RD: 10.1.1.1:12 (auto)
```

```
Import-RTs: 12:1
Export-RTs: 12:1
Per-EVI Label: none
State: Established
Encapsulation: vxlan
Vlan: 12 <<— VLAN Layer 2 VNI
Ethernet-Tag: 0
State: Established
Core If: Vlan99 <<— Interface handling IP VRF forwarding
Access If: Vlan12
NVE If: nve1
RMAC: 7035.0956.7edd <<— RMAC is the BIA of SVI 99 Core interface
Core Vlan: 99
L2 VNI: 12012
L3 VNI: 99999
VTEP IP: 10.255.1.1 <<— Local Tunnel endpoint IP address
MCAST IP: 239.12.12.12
VRF: vxlan <<— IP VRF for Layer 3 VPN
Pseudoports:
GigabitEthernet1/0/1 service instance 12
VTEP-2# show l2vpn evpn evi
EVI VLAN Ether Tag L2 VNI Multicast Pseudoport
----- ----- ---------- --------- ------------- ------------------
13 13 0 13013 239.13.13.13 Gi1/0/1:13
<<— See which EVPN instance maps to the VLAN. The VLAN
or EVPN instance values are not always the same
VTEP-2# show l2vpn evpn evi 13 detail
EVPN instance: 13 (VLAN Based)
RD: 10.2.2.2:13 (auto)
Import-RTs: 13:2
Export-RTs: 13:2
Per-EVI Label: none
State: Established
Encapsulation: vxlan
Vlan: 13 <<— VLAN Layer 2 VNI
Ethernet-Tag: 0
State: Established
Core If: Vlan99 <<— Interface handling IP VRF forwarding
Access If: Vlan13
NVE If: nve1
RMAC: 7486.0bc4.b75d <<— RMAC is the BIA of SVI 99 Core interface
Core Vlan: 99
L2 VNI: 13013
L3 VNI: 99999
VTEP IP: 10.255.2.1 <<— Local Tunnel endpoint IP address
MCAST IP: 239.13.13.13
VRF: vxlan <<— IP VRF for Layer 3 VPN
Pseudoports:
GigabitEthernet2/0/1 service instance 13
```

Verify that the Remote Layer 3 VNI Details are Learned on Each VTEP

The following examples show how to verify that the remote Layer 3 VNI details are learned on each VTEP:

```
VTEP-1# show nve peers
Interface VNI Type Peer-IP RMAC/Num_RTs eVNI state flags UP time
nve1 99999 L3CP 10.255.2.1 7486.0bc4.b75d 99999 UP A/M 1w1d
<<— Layer 3 Control Plane (L3CP), RMAC of Remote VTEP and Uptime of peer are displayed
VTEP-2# show nve peers
Interface VNI Type Peer-IP RMAC/Num_RTs eVNI state flags UP time
nve1 99999 L3CP 10.255.1.1 7035.0956.7edd 99999 UP A/M 21:27:36
<<— Layer 3 Control Plane (L3CP), RMAC of Remote VTEP and Uptime of peer are displayed
```

Verify that the Layer 3 VNI Tunnel Pseudoport is Installed into Layer 2 FIB in the Core VLAN

The following examples show how to verify that the Layer 3 VNI tunnel pseudoport is installed into Layer 2 FIB in the core VLAN:

```
VTEP-1# show l2fib bridge-domain 99 detail
<<— The Core VLAN can be obtained in the output of the
show l2vpn evpn evi <evpn-instance> detail command
Bridge Domain : 99
Reference Count : 8
Replication ports count : 0
Unicast Address table size : 1
IP Multicast Prefix table size : 3
Flood List Information :
Olist: 5112, Ports: 0
VxLAN Information :
Unicast Address table information :
7486.0bc4.b75d VXLAN_CP L:99999:10.255.1.1 R:99999:10.255.2.1
<<— Encapsulation Information to reach remote VTEP-2
IP Multicast Prefix table information :
Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5112, Ports: 0
Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5112, Ports: 0
Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5112, Ports: 0
VTEP-2# show l2fib bridge-domain 99 detail
<<— The Core VLAN can be obtained in the output of the
show l2vpn evpn evi <evpn-instance> detail command
Bridge Domain : 99
Reference Count : 8
Replication ports count : 0
Unicast Address table size : 1
IP Multicast Prefix table size : 3
Flood List Information :
Olist: 5111, Ports: 0
VxLAN Information :
Unicast Address table information :
7035.0956.7edd VXLAN_CP L:99999:10.255.2.1 R:99999:10.255.1.1
<<— Encapsulation Information to reach remote VTEP-2
IP Multicast Prefix table information :
Source: *, Group: 224.0.0.0/24, IIF: Null, Adjacency: Olist: 5111, Ports: 0
Source: *, Group: 224.0.1.39, IIF: Null, Adjacency: Olist: 5111, Ports: 0
Source: *, Group: 224.0.1.40, IIF: Null, Adjacency: Olist: 5111, Ports: 0
```

Verifying Traffic from a VXLAN Fabric to an IP Network Through a Border Leaf Switch Using Route Type 5

To verify the movement of traffic from a VXLAN fabric to an external IP network through a border leaf switch, perform these checks:

1. Check the Table Entries for BGP, EVPN, and VPNv4 Tables, on page 685
2. Check the Table Entries for BGP, EVPN, and VPNv4 Tables, on page 685
3. Confirm that Adjacency exists to Reach Tunnel Destination, on page 688

Check the Table Entries for BGP, EVPN, and VPNv4 Tables

The following examples show how to check the table entries for BGP, EVPN and VPNv4 tables:

```
VTEP-1# show bgp vpnv4 unicast vrf vxlan 10.9.9.9/32
```

```
   <<— To a remote IP address outside the VXLAN fabric
BGP routing table entry for 10.255.1.1:1:10.9.9.9/32, version 150
Paths: (1 available, best #1, table vxlan)    <<— VPNv4 VRF BGP table
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [5][10.255.1.11:1][0][32][10.9.9.9]/17 (global)
   <<— Learned from EVPN into VPNv4
    10.255.1.11 (metric 3) (via default) from 10.2.2.2 (10.2.2.2)
      Origin IGP, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 Router MAC:EC1D.8B55.F55D
      Originator: 10.255.1.11, Cluster list: 10.2.2.2
      Local vxlan vtep:
        vrf:vxlan, vni:99999
        local router mac:7035.0956.7EDD
        encap:8
        vtep-ip:10.255.1.1
        bdi:Vlan99
      Remote VxLAN:
        Topoid 0x2(vrf vxlan)
        Remote Router MAC:EC1D.8B55.F55D  <<— Border_Leaf_VTEP RMAC
        Encap 8
        Egress VNI 99999     <<— VNI associated with VRF
        RTEP 10.255.1.11     <<— Tunnel IP address
      rx pathid: 0, tx pathid: 0x0


VTEP-1# show bgp l2vpn evpn all route-type 5 0 10.9.9.9 32
   <<— This is sent as type 5 as there is no VNI at all for it to be mapped to
BGP routing table entry for [5][10.255.1.11:1][0][32][10.9.9.9]/17, version 650
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    10.255.1.11 (metric 3) (via default) from 10.2.2.2 (10.2.2.2)
   <<— Border_Leaf_VTEP Tunnel IP address
      Origin IGP, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 99999, MPLS VPN
Label 0
   <<— Using Layer 3 VNI 99999
      Extended Community: RT:10.255.1.11:1 ENCAP:8 Router MAC:EC1D.8B55.F55D
   <<— Route Target and RMAC of Border_Leaf_VTEP
      Originator: 10.255.1.11, Cluster list: 10.2.2.2
      rx pathid: 0, tx pathid: 0x0


Border_Leaf_VTEP# show bgp vpnv4 unicast vrf vxlan 10.12.12.12/32
   <<— To VXLAN Fabric IP address on VTEP-1
BGP routing table entry for 10.255.1.11:1:10.12.12.12/32, version 3092
Paths: (1 available, best #1, table vxlan)
Not advertised to any peer
Refresh Epoch 4
Local, imported path from [2][10.1.1.1:12][0][48][005F860210E7][32][10.12.12.12]/24 (global)

   <<— EVPN type-2 has been imported to VPNv4, from VTEP-1
  10.255.1.1 (metric 3) (via default) from 10.2.2.2 (10.2.2.2)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
    Extended Community: RT:10.255.1.11:1 ENCAP:8 Router MAC:7035.0956.7EDD
    Originator: 10.1.1.1, Cluster list: 10.2.2.2
    Local vxlan vtep:
      vrf:vxlan, vni:99999
      local router mac:EC1D.8B55.F55D
      encap:8
      vtep-ip:10.255.1.11
      bdi:Vlan99
```

```
    Remote VxLAN:
      Topoid 0x2(vrf vxlan)
      Remote Router MAC:7035.0956.7EDD  <<— VTEP-1 RMAC
       Encap 8
       Egress VNI 99999
       RTEP 10.255.1.1                  <<— VTEP-1 Tunnel IP address
       rx pathid: 0, tx pathid: 0x0

Border_Leaf_VTEP# show bgp l2vpn evpn all route-type 2 0 005F860210E7 10.12.12.12
   <<— Border_Leaf_VTEP still knows the type-2. This is still exchanged between the VTEPs
       even though the prefix has been imported to VPNv4
BGP routing table entry for [2][10.1.1.1:12][0][48][005F860210E7][32][10.12.12.12]/24,
version 3085
Paths: (1 available, best #1, table EVPN-BGP-Table)
Not advertised to any peer
Refresh Epoch 4
Local
10.255.1.1 (metric 3) (via default) from 10.2.2.2 (10.2.2.2)
  Origin incomplete, metric 0, localpref 100, valid, internal, best
  EVPN ESI: 00000000000000000000, Label1 12012, Label2 99999
   <<— Both Layer 2 VNI and Layer 3 VNI labels are seen in type-2,
       but only Layer 3 VNI 99999 is used, once imported to VPNv4
  Extended Community: RT:12:1 RT:10.255.1.1:1 ENCAP:8
    Router MAC:7035.0956.7EDD
  Originator: 10.1.1.1, Cluster list: 10.2.2.2
  rx pathid: 0, tx pathid: 0x0
```

Note To check if IP routes have been inserted into CEF table, run the **show ip route vrf** *vrf-name* command in privileged EXEC mode.

Confirm that Adjacency Exists to the VTEP Tunnel IP Address for the Host Device in IP VRF

The following examples show how to confirm that adjacency exists to the VTEP Tunnel IP address for the host device in IP VRF:

```
VTEP-1# show ip cef vrf vxlan 10.9.9.9/32 platform
10.9.9.9/32
  Platform adj-id: 0x1A, 0x0, tun_qos_dpidx:0   <<— Adjacency ID to remote IP address
```

```
VTEP-1# show platform software fed sw ac matm macTable vlan 99
VLAN   MAC                  Type  Seq#   EC_Bi  Flags  machandle          siHandle
      riHandle            diHandle             *a_time  *e_time  ports
-------------------------------------------------------------------------------------------
99     7035.0956.7edd      0x8002      0      0     64  0x7ffa48d61be8      0x7ffa48d630b8
      0x0                 0x5154                 0        0  Vlan99
99     7486.0bc4.b75d   0x1000001      0      0     64  0x7ffa48fb1bb8      0x7ffa48fac698
      0x7ffa48fab038      0x7ffa4838cc18            0        0  RLOC 10.255.2.1 adj_id
103
99     ec1d.8b55.f55d   0x1000001      0      0     64  0x7ffa48d065e8      0x7ffa48d01d08
      0x7ffa48c9a618      0x7ffa4838cc18            0        0  RLOC 10.255.1.11 adj_id
47
```

Confirm that Adjacency exists to Reach Tunnel Destination

The following example shows how to confirm that adjacency exists to reach Tunnel destination:

```
VTEP-1# show ip cef 10.255.1.11
10.255.1.11/32
  nexthop 10.1.1.6 GigabitEthernet1/0/1
```

Troubleshooting Tenant Routed Multicast

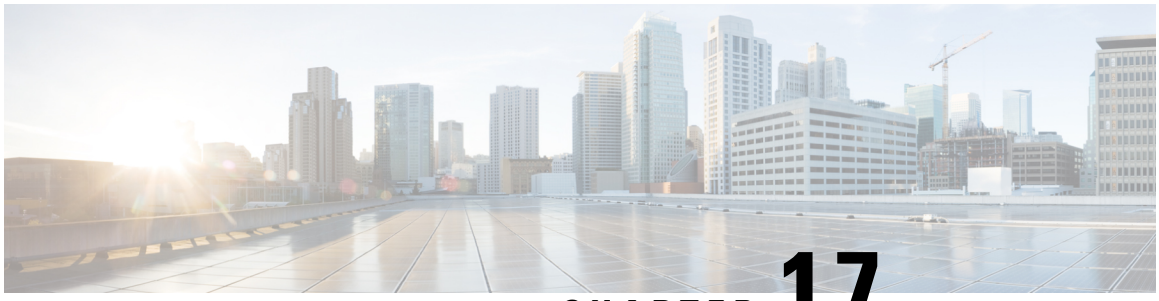
See Troubleshoot EVPN VxLAN TRM on Catalyst 9000 Switches document to learn how to troubleshoot issues with TRM in a BGP EVPN VXLAN fabric.



PART I

EVPN Multihoming: All-Active Mode

- All-Active Mode, on page 691
- Overlay: Routed, on page 703
- Overlay: Distributed Anycast Bridged, on page 727
- Overlay:Distributed Anycast Gateway Routed, on page 761



CHAPTER 17

All-Active Mode

The Ethernet VPN (EVPN) multihoming is RFC 7432, RFC 8365 based industry-standard Layer 2 multipath solution solving various traditional networking protocol challenges impacting scale, performance and resiliency. It is an advanced networking technology designed to provide high availability, Border Gateway Protocol (BGP) based Layer 2 loop-free network, and efficient load balancing in modern enterprise campus networks.

A pair of switches can connect to downstream Layer 2 or Layer 3 network device with local physical connection binding into single logical EtherChannel configured in Layer 2 trunk or access mode. With industry-standard non-blocking architecture, any downstream networking devices such as, Ethernet Switch, Wireless LAN Controller (WLC), firewalls, servers and hosts can be dual-homed.

Note All-active mode is enabled by default. To configure single-active mode, see Configuring Multihoming in a BGP EVPN VXLAN Fabric.

- EVPN multihoming key benefits, on page 691
- Restrictions for EVPN multihoming in all-active mode, on page 692
- EVPN multihoming technology overview, on page 693
- EVPN multihoming for non-fabric networks, on page 696
- EVPN multihoming in fabric networks, on page 697

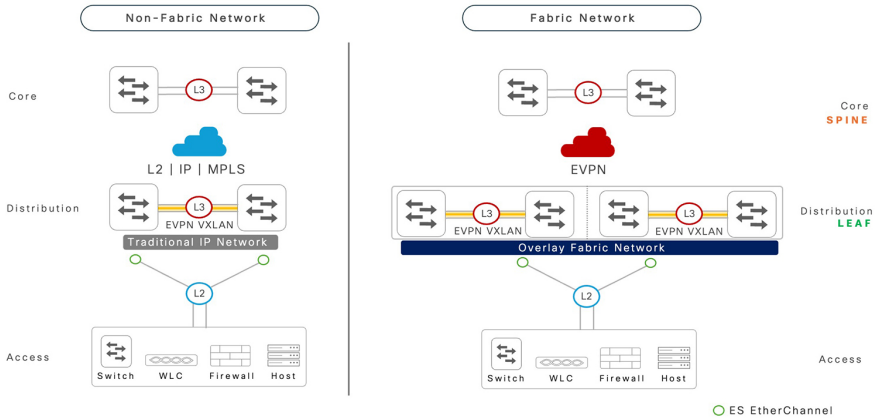
EVPN multihoming key benefits

EVPN multihoming technology was introduced on switches in Cisco IOS XE 26.2.1, delivering a highly flexible and extensible deployment solution for global enterprise customers.

The key technology architecture benefits provided by switches, combined with the advanced Cisco IOS XE software capabilities, enable the best of both worlds. Enterprise network administrators can implement EVPN multihoming within their existing traditional Layer 2 or Layer 3 network designs while laying a practical and gradual foundation for evolving towards modern, secure, fabric-based virtual networks.

The following illustration shows the EVPN multihoming solution for switches in two key flexible deployment models.

Figure 63: EVPN multihoming enterprise campus deployment model



With flexible deployment model support, EVPN multihoming provides significant technological benefits to support enterprise-grade networking requirements.

- **Industry-standard Layer 2 multipath**: RFC 7432 and RFC 8365-defined Layer 2 multipath technology with BGP control-plane eliminates the vendor proprietary complex solutions.
- **Flexible architecture**: Purpose-built solution supports a wide range of networking use-cases for traditional Layer 2 and Layer 3, to modernized secure campus fabrics enabling network-wide, two-tier virtual network segmentation and extensions.
- **Distributed planes**: A fully distributed architecture where each system maintains its own scalable and resilient control plane, independent management plane configuration, and software versions provide more resilience with synchronized data-plane between paired systems.
- **Increased performance**: Non-blocking Spanning Tree Protocol (STP), Layer 2 loop-free network combined with advanced load-sharing technique assist in increasing forwarding capacity and improving application experience.
- **Seamless integration**: Least disruptive integration in classic Layer 2 networks with a wide range of networking device types that support link-bundling capabilities. Flexible and seamless integration with existing core network system without a forklift architecture change.
- **Resilient**: Deterministic fault detection and recovery technique during various types of planned and unplanned failure conditions, which reduces the Mean Time to Repair (MTTR) and increases the network reliability with a higher Mean Time Between Failure (MTBF).

Restrictions for EVPN multihoming in all-active mode

This section lists the deployment restrictions for multihoming in all-active mode.

- EVPN multihoming for non-fabric and fabric networks with IPv6-only underlay network is not supported.
- EVPN multihoming in all-active mode is not supported on stacked switches.
- EVPN multihoming is not supported with Layer 2 leaf mode in centralized gateway mode fabric deployments.

- Anycast gateway (global or VRF-enabled) SVI with First Hop Redundancy Protocol (FHRP) protocol is not supported.
- Resilient Ethernet Protocol (REP) on Layer 2 Ethernet segment ports is not supported.
- Private VLAN (PVLAN) on Layer 2 Ethernet segment ports is not supported.
- DHCP snooping and ARP inspection on VLAN IDs mapped on logical Ethernet Segment EtherChannel trunk port is not supported.

EVPN multihoming technology overview

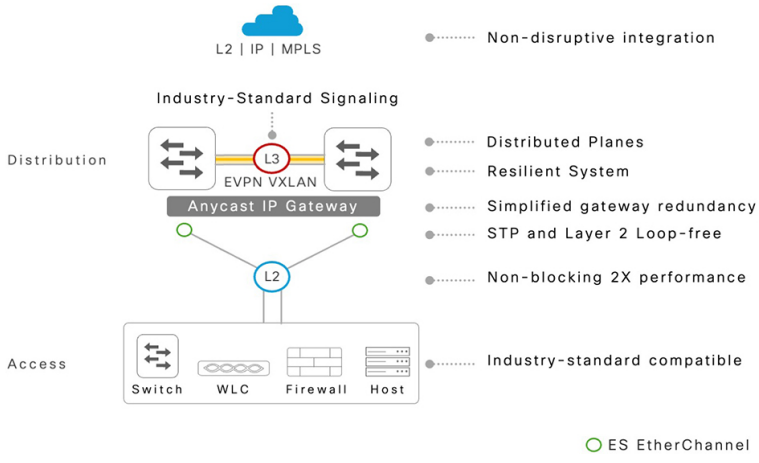
The legacy networking protocols in enterprise campus networks have been a challenge for IT organizations.

The challenges include difficulty in eliminating Spanning Tree Protocol (STP) that can lead to inefficient network topologies, restrictions in traditional Layer 2 designs that limit the network switching capacity, lack of deterministic reliability that makes it a challenge to support real-time mission-critical applications, insufficient support for mobility in wireless and legacy application environments, and increased complexity in network management and troubleshooting.

This section provides a brief overview of the advanced industry-standard EVPN multihoming technology for traditional Layer 2 and Layer 3 networks. Each subsection focuses on various technology components enabling loop-free, high-performance and resilient Layer 2 networking solutions.

The following illustration captures the key benefits of EVPN multihoming technology for simplified and resilient enterprise campus networks.

Figure 64: EVPN multihoming technology benefits



Refer to the *Terms and definitions* section for the description and purpose of acronyms.

Network system components

The EVPN multihoming technology comprises directly attached networking systems through traditional Layer 2 ports in trunk or access mode. The switches are required to be provisioned in EVPN multihoming mode to provide a loop-free Layer 2 multipath EtherChannel to any type of downstream network device.

Ethernet segment redundancy group

An Ethernet segment redundancy group is a physical grouping of network devices that form an EVPN multihoming network, enabling an all-active Layer 2 multipath solution.

The EVPN multihoming redundancy group consists of a single pair of switches, which provide physical connectivity and system-level redundancy. This architecture ensures continuous communication during both planned and unplanned failure events.

Ethernet segment switch

An Ethernet segment switch is a typical switch in core or distribution system configured to support the Layer 2 multihomed network with directly attached various types of Layer 2 systems.

The Ethernet Segment (ES) system can support multifunctional Layer 2 and Layer 3 networks for enterprise campus networks in classic and traditional non-fabric or modern fabric network roles. As a result, the ES switch is also known as a Provider Edge (PE) or Leaf or VTEP. Each ES switch operates and maintains fully distributed control, management, and data planes while solving legacy protocol challenges with industry-standard Layer 2 multipath solutions.

Ethernet segment switch platform and software

The EVPN multihoming technology operation between ES network devices—fixed or modular models, is no different from any of the other industry-standard networking protocols, like OSPF and BGP. A pair of ES switches can have different platform types, modules, interfaces, and so on. The ES pair switches can also have different Cisco IOS XE software versions to address stage upgrades and other conditions.

While the technology permits asymmetric platform and software versions on an ES pair, Cisco recommends a common platform and software version to deliver consistent performance and resiliency for non-disruptive business continuity during planned and unplanned failure conditions.

Ethernet segment client

The Layer 2 system is a directly attached single, dual, or multihomed connection to a pair of switches in an ES system. The Layer 2 ES client network can be of any device type but must follow the industry-standard Layer 2 networking technologies.

EVPN multihoming control plane

The EVPN multihoming technology is built on industry-standard, highly flexible, and proven BGP routing protocol. The BGP control plane replaces legacy STP for Layer 2 loop detection and prevention techniques.

The L2VPN EVPN address-family is a multiprotocol extension that enables a network-agnostic multihoming solution. The BGP protocol provides four key functions: discover remote neighbor and ES ID, real-time synchronization of network states, distributed forwarding rule to select local-bias rules for optimal performance, and resiliency for deterministic and efficient rapid fault-detection and recovery.

Ethernet segment

Ethernet segment (ES) is a pair of switches in the aggregation layer with a Layer 2 physical port that directly connects to an ES client device. It represents a single logical entity to enable loop-free, non-blocking all-active Layer 2 network connectivity.

An STP-Bridge Protocol Data Unit (BPDU) free Layer 2 network dynamically supports per-VLAN loop prevention while maximizing the network throughput to support accelerated application performance and resiliency.

Ethernet segment ID

An Ethernet segment ID is a 10-byte (00:01:01:01:01:01:01:01:01:01) identifier for each Layer 2 port that is connected to an ES client device. A common ES ID is assigned to a pair of ES systems on the Layer 2 physical port that connects to the same Layer 2 ES client device to enable all-active EVPN multihoming.

switches support 3 types of industry-standard ESI ID that is either auto-generated or IT-defined and manually configured.

Anycast gateway MAC and IP addresses

An SVI interface on a switch with anycast gateway is provisioned with a shared virtual IP address and an auto-derived shared virtual MAC address with **anycast-gateway mac auto** command in the global settings, on pair of EVPN multihoming switches. The unified, resolved ARP and ND addresses between the distributed IP gateway switches support optimal data load-balancing across all available bundled links and resilient networks during planned or unplanned failures for wired and wireless endpoints.

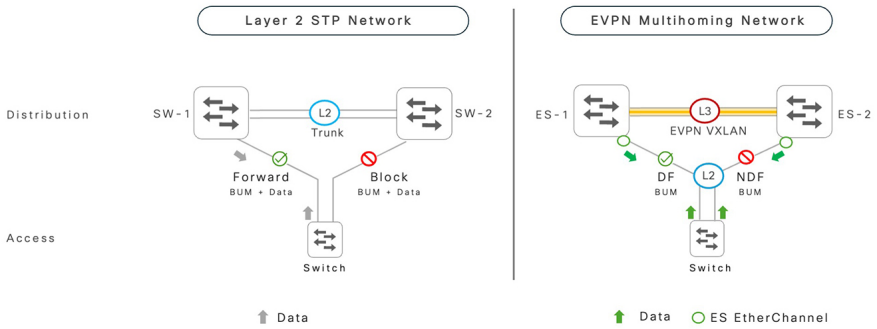
Designated forwarder and non-designated forwarder roles

Traditionally, in Layer 2 STP-enabled networks, loop detection is achieved through a blocking link to transmit broadcast and business application data traffic. As a result, the network operates inefficiently at reduced bandwidth capacity. Additionally, protocol-based fault detection and recovery increase the network convergence time, which impacts the reliability of mission-critical applications during faults.

In EVPN multihoming-based campus networks, the same result is obtained using a very different logic. Instead of blocking links, the EVPN distinctly decouples network traffic types between broadcast categories and business applications. On a per ES port and VLAN basis, the pair of switches dynamically assign specific roles to forward BUM (Broadcast, Unknown Unicast, and Multicast) traffic in active and passive modes. However, business application forwarding traffic is unblocked in active-active mode.

The switch elected to actively forward the BUM traffic on a shared ES EtherChannel interface is known as designated forwarder (DF). The switch elected to block the forwarding of the BUM traffic on a shared ES EtherChannel interface is known as non-designated forwarder (non-DF). The dynamic DF and non-DF role assignment is automatically derived based on the internal system modulo hash algorithm, which enables auto load balancing of the BUM traffic between a pair of switches connected to same downstream Layer 2 network access devices.

Figure 65: Designated forwarder and non-designated forwarder roles



Dynamic MAC and IP learning and synchronization

EVPN multihoming networks build and maintain the Layer 2 and Layer 3 network information using the control plane, between a pair of ES switches.

The upstream data towards the IP core network is hashed from Layer 2 network devices, the dynamic MAC, IPv4, and IPv6 host addresses synchronized using the BGP control plane in real time between both the ES switches.

The common MAC or IP forwarding tables enable high performance and fully-distributed local forwarding while pre-programming the inter-ES Layer 2 VXLAN tunnel to bridge the downstream Layer 2 network traffic as the last-resort interface to reroute the data plane rapidly, upon local path failure, without relying on data-plane flooding.

IGMP join and leave synchronization

The incoming Internet Group Management Protocol (IGMP) messages from multicast host receivers are locally processed by the connected ES switches. The IPv4 or IPv6 multicast group-to-IP membership information is synchronized between the ES switches by using an extended BGP control plane that supports consistent multicast state across the multihomed Ethernet segments.

The ES switch with a VLAN in the DF role transmits egress multicast traffic towards the receiver. The peer non-DF ES switch suppresses duplicate multicast frames to prevent loops and undesired multicast replication. This enables EVPN multihoming to support symmetric unicast or multicast application performance and resiliency during planned or unplanned failure events.

EVPN multihoming for non-fabric networks

EVPN multihoming in switches provides a flexible deployment solution for global enterprise customers to retain their traditional Layer 2 or Layer 3 networks in upstream core networks. Such network deployment option is also known as non-fabric network.

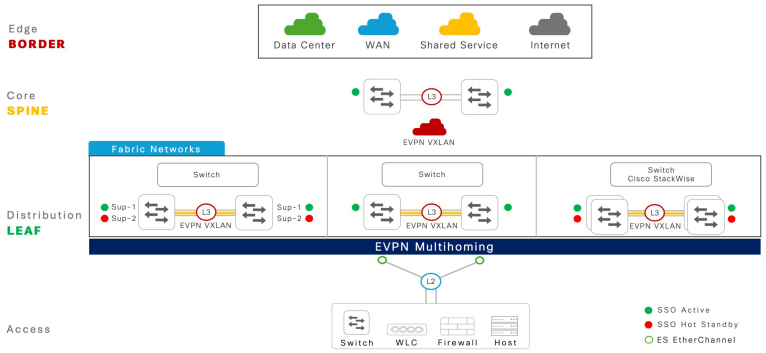
EVPN multihoming in fabric networks

Modern enterprise campuses require a secure fabric networking solution that seamlessly supports large-scale virtual and logical networks over a robust physical network infrastructure. EVPN multihoming-based networks offer a simplified and resilient architecture evolving from legacy STP protocol-based enterprise campus networks to a non-blocking Layer 2 fabric.

The EVPN VXLAN-enabled fabric core networks introduce new possibilities for EVPN multihomed Layer 2 networks by enabling secured wired and wireless segmented and extended networks that address critical technical and business requirements.

The following illustration displays EVPN multihoming in fabric networks

Figure 66: EVPN multihoming in fabric networks



Single unified BGP control plane

Enterprise campuses can be designed and deployed using a single unified BGP control plane that addresses the traditional Layer 2 and Layer 3 networking requirements at the access layer while supporting a modern fabric network at the core. The unified BGP control plane simplifies network operations by distinctly managing multi-domain routing and bridging functions, enabling seamless integration of the access and core networks.

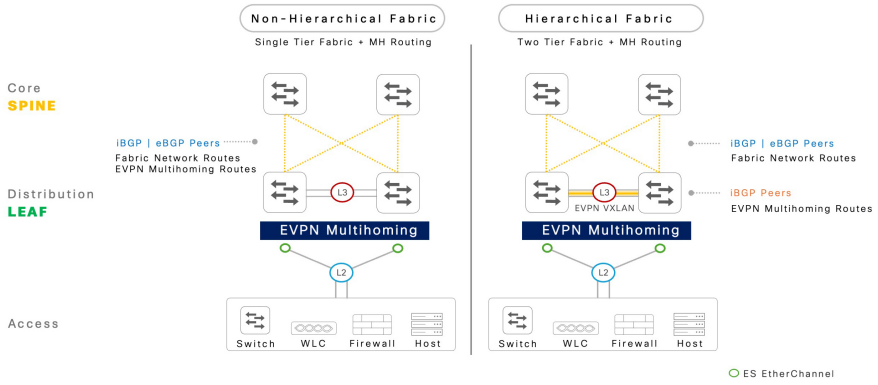
BGP EVPN multihoming-based networks represents a shift from the traditional flood-based networks to control-plane driven architecture that supports non-blocking and resilient enterprise campus networks. Host addresses—including the MAC, IPv4 or IPv6, multicast—are statefully discovered and synchronized across peer systems, enabling all-active data forwarding in BGP EVPN multihoming networks.

As the BGP control plane role expands into traditional networks, processing demands may increase due to additional BGP prefix tables supporting EVPN multihoming auto-generated prefixes.

switches support both standard non-hierarchical and hierarchical BGP peering mode between leaf and spine devices, facilitating the exchange of fabric and EVPN multihoming network prefixes.

To address scalability and performance requirements, switches can implement a two-tier hierarchical control-plane network in large deployments.

Figure 67: BGP peering alternatives for fabric networks



Non-hierarchical fabric networks

The general non-hierarchical BGP EVPN VXLAN network implementation follows a standard single tier BGP peering model where all fabric devices peer directly with each other to dynamically exchange routes and build the VXLAN forwarding tunnels.

The EVPN multihoming-enabled non-hierarchical fabric network implementation follows a standard iBGP or eBGP peering model between the spine and EVPN multihoming-enabled leaf or border system.

In large scale enterprise campus network environments, spine layer devices require additional resources such as processing power and memory to handle standard fabric host information, such as MAC addresses, MAC/IP bindings and network prefixes (IPv4/IPv6) along with EVPN multihoming auto-generated routes received from each leaf or border system.

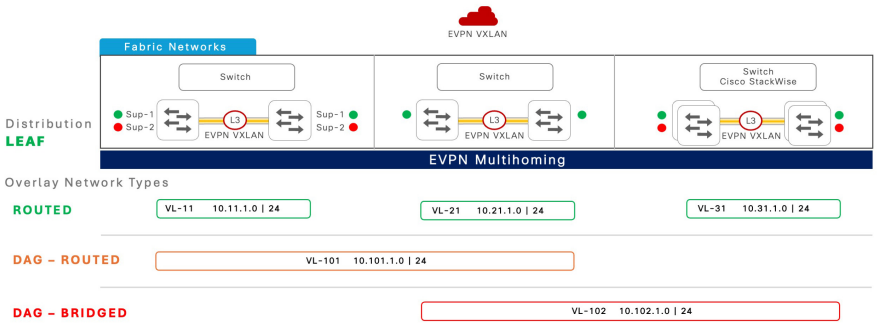
EVPN fabric overlay network types overview

BGP EVPN VXLAN provides a flexible, virtual networking solution that addresses key networking use cases by combining a wide range of overlay architectures.

Such business-driven network architectures enable multidimensional benefits for scalable and secure segmented overlay networks in EVPN multihoming networks. Network administrators can build overlay networks with a “route first” mindset that can conditionally extend IP or VLAN networks between targeted Ethernet segment network devices to meet application requirements.

The following figure illustrates three overlay network types in EVPN multihoming networks.

Figure 68: BGP EVPN fabric overlay network types



Network administrators decide on the type of overlay networks to deploy across the fabric core based on the specific application requirements to achieve better scale and resiliency in EVPN multihoming networks.

switches provide a flexible solution that supports coexistence of all overlay network types within a single system. Based on technical requirements, each VLAN from an Ethernet segment EtherChannel on switches can be configured to support IP-routed network segmentation, conditionally stretch IP subnets, or bridge VLANs using Layer 2 flood mechanisms selectively between VLANs.

Table 61: Comparison of EVPN overlay network types

	EVPN overlay network types		
	Routed	**DAG routed**	**DAG bridged**
Function	IP routing in core and Layer 3 segmented overlay.	Flood-free IP subnet stretch and Layer 3 segmented overlay.	Layer 2 flood stretch and Layer 3 segmented overlay.
Use case	Wired or wireless data. IT or OT endpoints	Wired or wireless data. IT or OT endpoints	Wired or wireless data. IT or OT endpoints. Non-IP endpoints
VLAN/subnet	One VLAN or subnet per distribution block	Stretched IP subnets between targeted distribution blocks	Stretched VLANs or subnets between targeted distribution blocks
Layer 2 flood boundary	Within the local Layer 2 network	Within the local Layer 2 network	Within the local Layer 2 network and across the fabric core
IP gateway	Anycast gateway per distribution block	Anycast gateway between targeted distribution blocks	Anycast gateway between targeted distribution blocks
IP subnet stretch	Not applicable	Yes—flood-free	Yes—flood-based
Seamless distributed wireless mobility	Within the local distribution block	Within the local distribution block and across targeted destination blocks	Within the local distribution block and across targeted destination blocks

	EVPN overlay network types		
	Routed	**DAG routed**	**DAG bridged**
Non-IP/silent host support	Within the local distribution block	Within the local distribution block	Within the local distribution block and across targeted destination blocks
Recommendation	Best scalable solution	Extends IP subnet selectively; use only if required	Extends Layer 2 flood selectively; use only if required

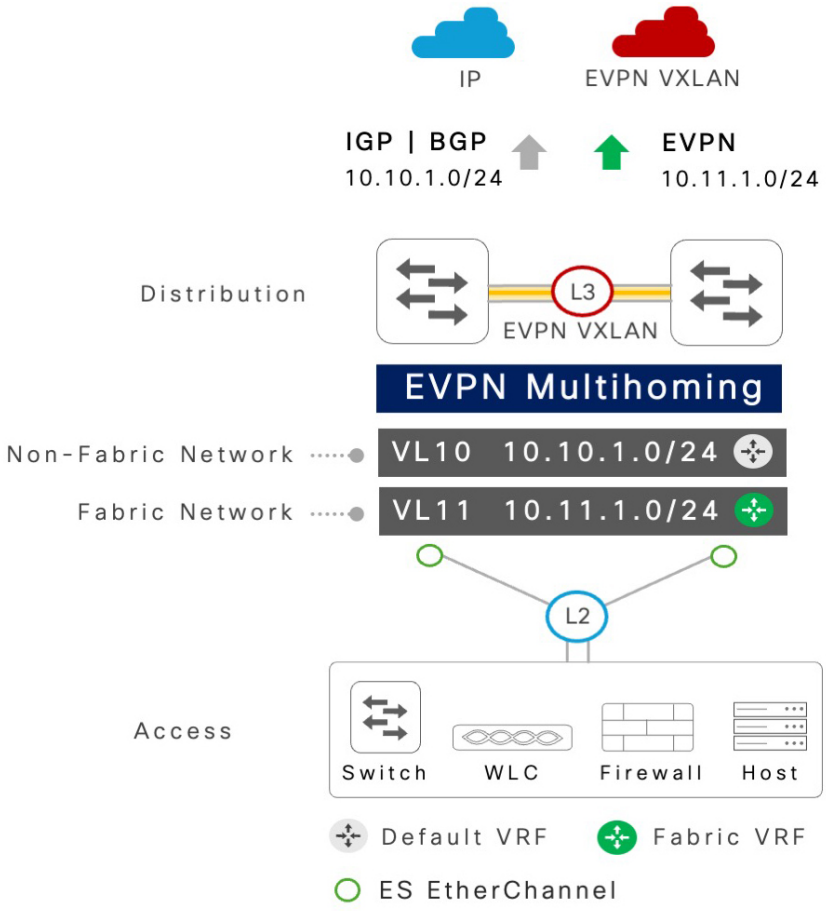
EVPN fabric and non-fabric network co-existence

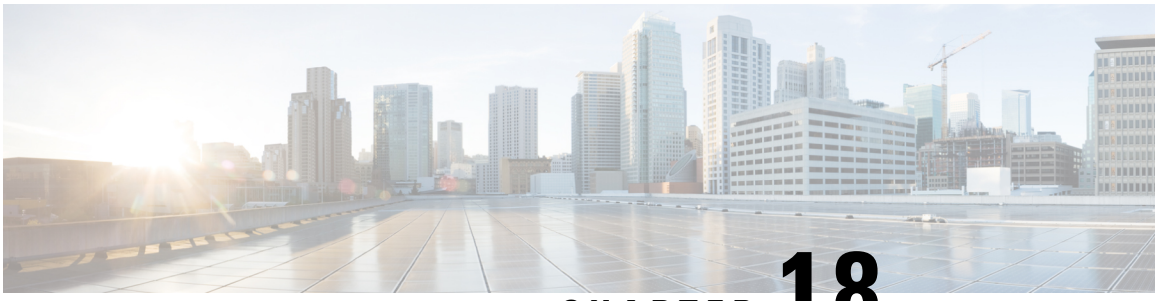
Enterprise campuses require flexible networking solutions that support both modern secure fabric networks and backward-compatible traditional Layer 2 or Layer 3 networks. switches with EVPN multihoming-enabled EtherChannel trunk interfaces provide this flexibility by bundling multiple VLANs, each delivering distinct networking services.

Network administrators can program some VLANs to operate traditional IP-based networks in the underlay, while mapping other VLANs to IP VRFs or MAC VRFs to route and bridge over the VXLAN fabric.

The following illustration shows a network deployment scenario with both fabric and non-fabric VLANs on EVPN multihoming-enabled networks.

Figure 69: EVPN multihoming: fabric and non-fabric network co-existence





CHAPTER 18

Overlay: Routed

Routed overlays enable the separation of underlay and overlay routing, allowing for flexible and scalable network designs.

- Routed overlay network overview, on page 703
- Understanding routed overlay networks, on page 704
- Configure routed overlay networks, on page 705
- Verify the routed overlay network configuration, on page 713
- Reference routed overlay network configuration, on page 715

Routed overlay network overview

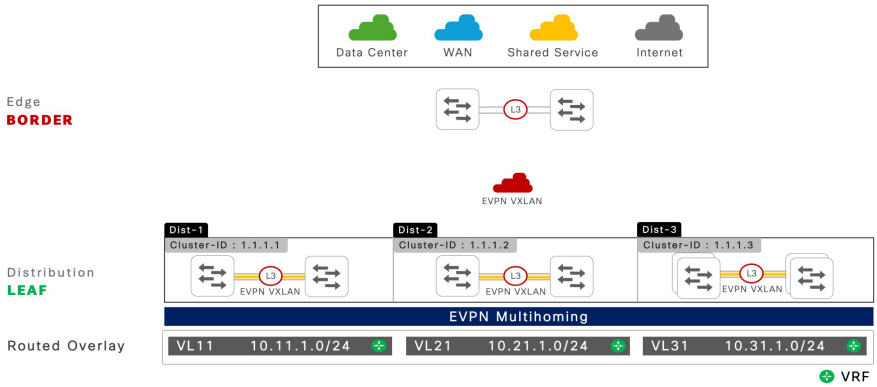
In multilayer campus networks, the traditional IP addressing plan follows a VLAN-to-subnet approach on a per distribution block basis.

The networking requirements and use cases for business applications typically remain within a single Layer 2 or Layer 3 network. The unique IPv4 or IPv6 subnet from each distribution network can be IP-routed in the default routing space using traditional Interior Gateway Protocols (IGPs) or BGP protocols in the underlay network, or the subnet can be logically segmented and routed in the overlay VXLAN fabric network.

The BGP EVPN VXLAN fabric routed overlay is highly relevant when deploying switches in the leaf role at the campus distribution layer. The unified and structured IP addressing plan between the underlay and overlay networks provides a congruent design, deployment, and operational experience across the enterprise network.

The following figure shows a reference EVPN multihoming network with a routed overlay deployment scenario.

Figure 70: EVPN multihoming with routed overlay fabric network

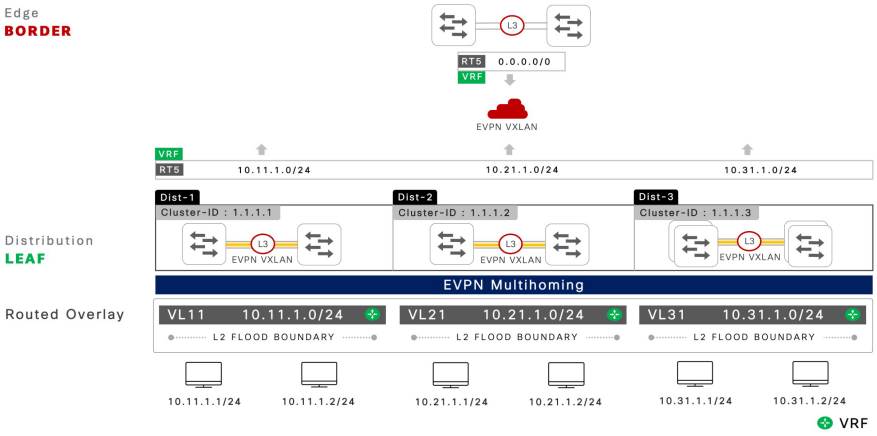


Understanding routed overlay networks

The fundamentals of Layer 2 and Layer 3 networking with routed overlay in BGP EVPN VXLAN-based fabric networks remain unchanged from the traditional campus network deployment models. This section provides the key functional and operational components for building scalable BGP EVPN VXLAN fabric networks with EVPN multihoming.

The following figure shows the characteristics of an EVPN multihoming network with routed overlay fabric.

Figure 71: Key fabric network characteristics of a routed overlay network



Hierarchical BGP control plane

In large-scale campus networks, the recommended deployment approach for controlled route management between EVPN multihoming peers and spines is a two-tier hierarchical BGP control plane. This design enables each BGP peering layer to manage domain-specific EVPN routes to support downstream non-blocking and all-active Layer 2 multipath network, and manage the advertisement of fabric network prefixes to enable secure global network access connectivity.

EVPN multihoming

The Layer 2 network with EVPN multihoming operation functions independently from the EVPN fabric core network. The pair of switches at each distribution layer form and manage EVPN multihoming within their respective redundancy group. These switches use a unified Layer 2 Ethernet Segment (ES) EtherChannel connection to downstream network devices. The EtherChannel connection can carry VLANs that are either mapped to routed overlay and advertised within the fabric core or continue to operate within the traditional underlay IP network.

Layer 2 broadcast network boundary

The EVPN fabric routed overlay networks provide a unified Layer 2 or Layer 3 network boundary similar to the traditional underlay campus networks. The Layer 2 bridge-domain and blast radius are contained between the Layer 2 access layer and directly attached EVPN multihoming-enabled distribution layer systems.

The pair of switches in the distribution layer network dynamically builds a Layer 2 VXLAN tunnel using ingress replication mode to extend Broadcast, Unknown Unicast and Multicast (BUM) traffic. With inbuilt loop detection and fast convergence techniques, the EVPN multihoming-enabled networks support secure and scalable Layer 2 networks over traditional STP-based networks.

Structured IP subnet and gateway redundancy

The Distributed Anycast Gateway (DAG) routed overlay networks follow the standard enterprise campus design principles with a VLAN or subnet per distribution block. Network administrators can define a structured IPv4 or IPv6 addressing plan to build scalable overlay networks.

A pair of switches operating in EVPN multihoming mode provides IP gateway redundancy by implementing a built-in gateway function for both IPv4 or IPv6 in each distribution block.

Spine-route policy

In large scale EVPN multihoming enabled fabric networks, controlled route management is enforced through route policies on leaf switches that are attached to spine switches. The iBGP or eBGP prefix advertisement to spine switches in the fabric core limits the advertising of IPv4 or IPv6 network prefixes, specifically EVPN route type 5 prefixes, while retaining unfiltered routing with iBGP peering sessions between leaf switches and remote ES switches. This design supports a non-blocking Layer 2 multipath EtherChannel that ensures scalability and efficient route distribution.



Note EVPN multihoming-enabled networks support additional overlay network types to stretch the IP subnets or Layer 2 flood-domains beyond a single distribution block by using the DAG technology. Routed overlay networks are recommended for better scale, performance, and resiliency.

Configure routed overlay networks

This section lists the tasks to configure a routed overlay network.

Prerequisites to configure routed overlay networks

A reliable underlay IP network and a Layer 2 EVPN multihoming network are prerequisites for configuring a BGP EVPN VXLAN-enabled fabric core network.

Configure IP VRFs

Perform this task to configure a logically segmented virtual network with an IP VRF to exchange IPv4 or IPv6 network prefixes across all the VTEPs in the fabric core network.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition green`	Configures a new Virtual Routing and Forwarding (VRF) instance and enters VRF configuration mode.
Step 4	**rd** *route distinguisher* **Example:** `Device(config-vrf)# rd 10.200.255.101:101`	Configures a route distinguisher with distinct values across all fabric devices.
Step 5	**address-family ipv4** [**multicast** | **unicast**] **Example:** `Device(config-vrf)# address-family ipv4 unicast`	Enters IPv4 unicast address-family configuration mode.
Step 6	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **Example:** `Device(config-vrf-af)# route-target 65101:101`	Configures overlay IPv4 routing import or export policies that match extended community values. • The default **both** keyword will auto-generate the import and export command-line with the configured community value.
Step 7	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **stitching** **Example:** `Device(config-vrf-af)# route-target 65101:101 stitching`	Configures overlay IPv4 routing import or export policies that match the extended community values to enable extended Network Layer Reachability Information (NLRI) types for BGP EVPN VXLAN enabled networks. The default value is **both**.

	Command or Action	Purpose
Step 8	**exit** **Example:** `Device(config-vrf-af)# exit`	Exits IPv4 unicast address-family configuration mode and returns to VRF configuration mode.
Step 9	**address-family ipv6** [**multicast** | **unicast**] **Example:** `Device(config-vrf)# address-family ipv6 unicast`	Enters IPv6 unicast address-family configuration mode.
Step 10	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **Example:** `Device(config-vrf-af)# route-target 65101:101`	Configures overlay IPv6 routing import and export policies that match the extended community value. • The default **both** keyword will auto-generate the import and export command-line with the configured community value.
Step 11	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **stitching** **Example:** `Device(config-vrf-af)# route-target 65101:101 stitching`	Configures the overlay IPv6 routing import and export policies that match the extended community values to enable extended NLRI types for BGP EVPN VXLAN enabled networks. The default value is **both**.
Step 12	**end** **Example:** `Device(config-vrf-af)# end`	Exits IPv6 unicast address-family configuration mode and returns to privileged EXEC mode.

Configure core VLAN and SVI interfaces

Perform this task to configure core VLAN and SVI interfaces for each IP VRF to enable VXLAN encapsulation and decapsulation functions between downstream non-VXLAN network and remote VXLAN-enabled VTEPs.

Note Configure a Network Virtualization Endpoint (NVE) interface and bind the L3VNI ID to a logical interface.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan** *id* **Example:** `Device(config)# vlan 101`	Configures a dedicated VLAN ID for IP VRF core routing and enters VLAN configuration mode. • This VLAN must be pruned from Layer 2 trunk ports and cannot be used for any other purpose.
Step 4	**vlan configuration** *id* **Example:** `Device(config-vlan)# vlan configuration 101`	Configures a core VLAN to assign the Layer 3 VXLAN network (VNI) ID.
Step 5	**member vni** *l3vni-id* **Example:** `Device(config-vlan)# member vni 11101`	Configures a unique Layer 3 VNI ID for an IP VRF. • For example, VRF Green will have VLAN 101 as the dedicated VLAN and Layer 3 VNI, and this VNI ID must not be shared with other IP VRFs on the same system.
Step 6	**exit** **Example:** `Device(config-vlan)# exit`	Exits VLAN configuration mode and returns to global configuration mode.
Step 7	**interface vlan** *id* **Example:** `Device(config-vlan)# interface vlan 101`	Configures a core VLAN SVI interface and assigns it to an IP VRF and enters interface configuration mode. • The core VLAN SVI MAC address is used to support routed MAC address for Layer 3 overlay networks.
Step 8	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Maps IP VRF to a core VLAN SVI interface.
Step 9	**ip unnumbered** *interface-type interface-number* **Example:** `Device(config-if)# ip unnumbered Loopback0`	Configures unnumbered IPv4 address from an underlay logical interface.

	Command or Action	Purpose
Step 10	**no autostate** **Example:** `Device(config-if)# no autostate`	Disables autostate on the interface. • In EVPN deployments, a VLAN used for a core-facing SVI should not be allowed in any trunk ports. For a core-facing SVI to function properly, the **no autostate** command must be configured under the SVI.
Step 11	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 12	**interface nve** *id* **Example:** `Device(config)# interface nve 1`	Configures a virtual NVE interface to bind one or more L2VNIs and L3VNIs to enable the system-wide VXLAN forwarding function • Enters NVE interface configuration mode.
Step 13	**source-interface** *id* **Example:** `Device(config-if-nve)# source-interface Loopback0`	Configures the loopback interface as the source address to enable data communication over a VXLAN tunnel.
Step 14	**host-reachability protocol bgp** **Example:** `Device(config-if-nve)# host-reachability protocol bgp`	Configures BGP as the host reachability control plane protocol on the interface.
Step 15	**member vni** *l3vni-id* **vrf** *vrf-name* **Example:** `Device(config-if-nve)# member vni 11101 vrf green`	Binds the L3VNI and IP VRF under the NVE interface.
Step 16	**end** **Example:** `Device(config-if-nve)# end`	Exits NVE interface configuration mode and returns to privileged EXEC mode.

Configure access SVI interfaces

This section provides step-by-step instructions to configure two Ethernet Segment (ES) systems as a single logical IP gateway with the network edge SVI interface sharing a common IP address.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *id* **Example:** `Device(config)# interface vlan 11`	Configures a Layer 3 access SVI interface ID and enters interface configuration mode.
Step 4	**ip address** *ip-address mask* **Example:** `Device(config-if)# ip address 10.11.1.254 255.255.255.0`	Assigns an IPv4 gateway address and mask for the selected VLAN. Anycast gateway is auto-enabled with a common IP address between the ES systems.
Step 5	**ipv6 address** *ipv6-address mask* **Example:** `Device(config-if)# ipv6 address 2001:10:11:1::f/64`	Assigns an IPv6 gateway address and prefix for the selected VLAN. Anycast gateway is auto-enabled with a common address between the ES systems.
Step 6	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and returns to privileged EXEC mode.

Configure spine EVPN route policies

This section provides step-by-step instructions to configure two switches that act as a pair of ES systems with an outbound spine EVPN route policy. The route policy controls the advertisement of EVPN routes by the ES systems for a scalable routed overlay network, independent of iBGP or eBGP spine peering session type.

The route policy limits the advertisement of route type 5 (RT-5) prefixes to spine switches for network-wide route propagation. iBGP peering with EVPN multihoming system maintains unfiltered EVPN route exchange to build and maintain a non-blocking Layer 2 multipath network.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode • Enter your password if prompted.

	Command or Action	Purpose
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**route-map** *name* [**permit** | **deny**] *sequence-number* **Example:** `Device(config)# route-map SPINE-ROUTE-POLICY-OUT permit 10`	Creates a route-map rule with permit or deny match criteria for filtering or modifying routing information. The optional *sequence-number* argument indicates the order of route-map rule processing. • Enters route-map configuration mode.
Step 4	**description** *description* **Example:** `Device(config-route-map)# description ROUTED OVERLAY – VL11`	(Optional) Configures a description for the route-map sequence.
Step 5	**match evpn route-type** [**1-8** | **2-mac-ip** | **2-mac-only**] **Example:** `Device(config-route-map)# match evpn route-type 5`	Configures an EVPN route policy sequence to permit the advertisement of local IPv4 or IPv6 RT-5 network prefixes to any IP VRF.
Step 6	**exit** **Example:** `Device(config-route-map)# exit`	Exits route-map configuration mode and returns to global configuration mode.
Step 7	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 65101`	Configures a BGP Autonomous System Number (ASN) using 2-bytes or 4-bytes in asplain and asdot formats and enters router configuration mode.
Step 8	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Enters BGP L2VPN address-family configuration mode.
Step 9	**neighbor** *spine-loopback-address* **route-map** *name* **out** **Example:** `Device(config-router-af)# neighbor 10.200.255.3 route-map SPINE-ROUTE-POLICY-OUT out`	Applies the outbound route-map policy to each spine peer to limit advertising the RT-5 prefix.
Step 10	**end** **Example:** `Device(config-router-af)# end`	Exits BGP L2VPN address-family configuration mode and returns to privileged EXEC mode.

Configure IP VRF address family routing

Perform this task to advertise network prefixes in IPv4 and IPv6 address family in overlay networks.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 65101`	Configures a BGP ASN using 2-bytes or 4-bytes in asplain and asdot formats and enters router configuration mode.
Step 4	**address-family** [**ipv4** | **ipv6**] **vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv4 vrf green`	Enables IPv4 address-family for a VRF instance to enable IPv4 network prefix routing in overlay networks.
Step 5	**advertise l2vpn evpn** **Example:** `Device(config-router-af)# advertise l2vpn evpn`	Enables the advertising of the IPv4 network prefix routing between an IPv4 VRF instance and an EVPN fabric core network.
Step 6	**redistribute connected** [**metric** | **route-map** *name*] **Example:** `Device(config-router-af)# redistribute connected`	Configures an IPv4 address family to advertise all locally connected IPv4 network prefixes in the EVPN fabric network. • You can also advertise the network prefix using the **network** command.
Step 7	**exit** **Example:** `Device(config-router-af)# exit`	Exits IPv4 address-family configuration mode and returns to router configuration mode.
Step 8	**address-family** [**ipv4** | **ipv6**] **vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv6 vrf green`	Enables IPv6 address family for a VRF instance to enable IPv6 network prefix routing in overlay networks.

	Command or Action	Purpose
Step 9	**advertise l2vpn evpn** **Example:** `Device(config-router-af)# advertise l2vpn evpn`	Enables IPv6 network prefix routing between an IPv6 VRF instance and the EVPN fabric core network.
Step 10	**redistribute connected** [**metric** | **route-map**] *name* **Example:** `Device(config-router-af)# redistribute connected`	Configures an IPv6 address family to advertise all locally connected IPv6 network prefixes in the EVPN fabric network. • You also can use the **network** command to advertise the network prefix.
Step 11	**end** **Example:** `Device(config-router-af)# end`	Exits IPv6 address family configuration mode and returns to privileged EXEC mode.

Verify the routed overlay network configuration

This section lists the commands and the output of these commands to verify the routed overlay network configuration and the states of the VTEPs.

The command output may be truncated to focus on critical information useful for Day Two operations and troubleshooting.

IP VRF networks: Verifies the locally configured IP VRFs and the associated physical or logical interface bindings to a virtual network on VTEPs, including leaf and border devices.

The interface bindings displayed in this command output include the network edge VLANs connecting to endpoints, network devices along with a single core VLAN (VLAN 101) of an IP VRF.

```
ES-1# show vrf green

  Name          Default RD            Protocols Interfaces
  green         10.200.255.103:101     ipv4  Vl11
                                                 Vl101

BORDER-1# show vrf green

  Name           Default RD            Protocols Interfaces
  green          10.200.255.1:101 ipv4  Vl101
                                                  Vl1101
```

BGP L2VPN EVPN neighbors: Verifies that the two-tier hierarchical iBGP session between a pair of switches in EVPN multihoming mode and the iBGP peering to a pair of spine systems are operational.

The following output displays iBGP peering between a pair of spine switches, 10.100.255.3 and 10.100.255.4 and direct iBGP peering between ES-1 local 10.100.255.101 and ES-2 10.100.255.102 switches are operational and in EVPN multihoming mode.

```
ES-1# show bgp l2vpn evpn all summary
```

```
BGP router identifier 10.100.255.101, local AS number 65101
<snip>

Neighbor           V        AS      MsgRcvd MsgSent   TblVer  InQ     OutQ    Up/Down
State/PfxRcd
10.100.255.3       4        65101      18     20      104      0        0     00:04:35
    2
10.100.255.4       4        65101      23     25      106      0        0     00:05:19
    2
10.100.255.102     4        65101      51     65      104      0        0     00:04:26
    28
```

Spine policy: Verifies that a VTEP in leaf or border roles implemented in EVPN multihoming mode is configured with a route map and has an outbound policy applied to each spine, such as 10.200.255.3. The spine policy configuration on the border switch is optional and not required if the switch is not implemented in EVPN multihoming mode to connect to external network devices, such as firewalls.

```
ES-1# show bgp l2vpn evpn neighbor 10.200.255.3 policy

 Neighbor: 10.200.255.3, Address-Family: L2VPN E-VPN
 Locally configured policies:
  send-community extended
 Inherited polices:
  route-map EVPN-SPINE-ROUTE-POLICY-OUT out
```

L2VPN advertised routes: Verifies that a VTEP in leaf or border roles implemented in EVPN multihoming mode is advertising EVPN prefixes to the spine (10.200.255.3) based on the applied route-map policy.

The following output shows that the RT-5 EVPN prefix is advertised, and no additional EVPN route-types are advertised to the spine switches.

```
ES-1# show bgp l2vpn evpn neighbor 10.200.255.3 advertised-route

BGP table version is 7179, local router ID is 10.200.255.101
<snip>
     Network          Next Hop            Metric  LocPrf  Weight  Path
Route Distinguisher: 10.200.255.101:101 (default for vrf green)
 *>   [5][10.200.255.101:101][0][24][10.1.11.0]/17
                       0.0.0.0                  0           32768  ?
 *>   [5][10.200.255.101:101][0][24][10.1.12.0]/17
                       0.0.0.0                  0           32768  ?
 *>   [5][10.200.255.101:101][0][24][10.1.13.0]/17
                       0.0.0.0                  0           32768  ?


Border-1# show bgp l2vpn evpn neighbors 10.200.255.3 advertised-routes

BGP table version is 2895, local router ID is 10.200.255.1
<snip>
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 10.100.253.23:1101 (default for vrf green)
 *>   [5][10.100.253.23:1101][0][0][0.0.0.0]/17
                       21.1.1.1                 0              0  65001 ?
 *>   [5][10.100.253.23:1101][0][16][111.1.0.0]/17
                       21.1.1.1                 0              0  65001 ?
```

VRF IP routing table: Verifies that a VTEP in leaf or border roles implemented in EVPN multihoming mode is advertising the EVPN prefixes to the spine, 10.200.255.3 based on applied route-map policy.

The following outputs from leaf and border roles confirm that only RT-5 EVPN prefix is advertised, and no additional EVPN route types are advertised to the spine switches.

```
ES-1# show ip route vrf green bgp

Routing Table: green
<snip>

Gateway of last resort is 10.200.255.2 to network 0.0.0.0
B*    0.0.0.0/0    [200/0] via 10.200.255.1, 00:00:32, Vlan101
                   [200/0] via 10.200.255.2, 00:00:32, Vlan101


Border-1# show ip route vrf green bgp
Routing Table: green
<snip>
Gateway of last resort is 21.1.1.1 to network 0.0.0.0

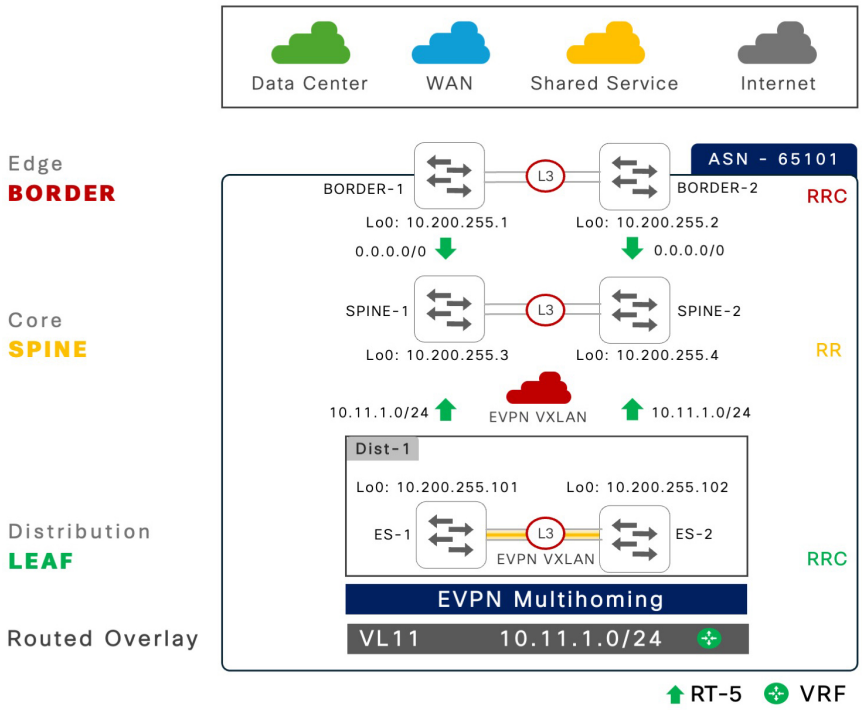
B*    0.0.0.0/0 [20/0] via 21.1.1.1, 6w0d
      10.0.0.0/24 is subnetted, 95 subnets
B     10.11.1.0    [200/0] via 10.200.255.104, 00:04:25, Vlan1101
                   [200/0] via 10.200.255.103, 00:04:25, Vlan1101
```

Reference routed overlay network configuration

This section provides a reference network design to implement an end-to-end routed overlay fabric network.

The following figure illustrates an EVPN multihoming network with hierarchical BGP peering that involves the spine and leaf architecture.

Figure 72: EVPN multihoming: Routed overlay fabric network design



Layer 2 EVPN multihoming

This step-by-step configuration shows how to build a non-blocking all-active Layer 2 network with EVPN multihoming technology on a pair of switches, starting from a Layer 2 campus network.

Step	ES-1	ES-2
1: Inter-ES Layer 3 EtherChannel	! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-2 no switchport ip address 10.0.0.0 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 101 ip ospf cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !	! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-1 no switchport ip address 10.0.0.1 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 101 ip ospf cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !

Step	ES-1	ES-2
2: IGP routing and core interface	! router ospf 100 router-id 10.200.255.101 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low area 101 stub no-summary passive-interface default no passive-interface Port-Channel 128 no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 ! interface Loopback 0 ip address 10.100.255.101 255.255.255.255 ip ospf 100 area 0 !	! router ospf 100 router-id 10.200.255.102 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low area 101 stub no-summary passive-interface default no passive-interface Port-Channel 128 no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 ! interface Loopback 0 ip address 10.100.255.102 255.255.255.255 ip ospf 100 area 0 !
3: iBGP routing	! router bgp 65101 template peer-policy ES-PEER-POLICY send-community both ! template peer-session ES-PEER-SESSION-POLICY remote-as 65101 description EVPN-MH-DIST-2-PEER update-source Loopback0 fall-over host-route ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast neighbor 10.100.255.102 inherit peer-session ES-PEER-SESSION-POLICY ! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.100.255.102 activate neighbor 10.100.255.102 send-community both neighbor 10.100.255.102 inherit peer-policy ES-PEER-POLICY !	! router bgp 65101 template peer-policy ES-PEER-POLICY send-community both ! template peer-session ES-PEER-SESSION-POLICY remote-as 65101 description EVPN-MH-DIST-1-PEER update-source Loopback0 fall-over host-route ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast neighbor 10.100.255.101 inherit peer-session ES-PEER-SESSION-POLICY ! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.100.255.101 activate neighbor 10.100.255.101 send-community both neighbor 10.100.255.101 inherit peer-policy ES-PEER-POLICY !

Step	ES-1	ES-2
4: Global L2VPN	! l2vpn evpn advertise mac disable anycast-gateway mac auto multicast advertise sync-only multihoming aliasing disable multihoming peering adjacent replication-type ingress router-id Loopback 0 !	! l2vpn evpn advertise mac disable anycast-gateway mac auto multicast advertise sync-only multihoming aliasing disable multihoming peering adjacent replication-type ingress router-id Loopback 0 !
5: Routed VLAN and MAC VRF	! vlan 11 name ROUTED_DATA_VLAN ! l2vpn evpn instance 11 vlan-based encapsulation vxlan ! vlan configuration 11 member evpn-instance 11 vni 11011 ! interface nve 1 source-interface Loopback 0 host-reachability protocol bgp member vni 11011 ingress-replication !	! vlan 11 name ROUTED_DATA_VLAN ! l2vpn evpn instance 11 vlan-based encapsulation vxlan ! vlan configuration 11 member evpn-instance 11 vni 11011 ! interface nve 1 source-interface Loopback 0 host-reachability protocol bgp member vni 11011 ingress-replication !
6: ES EtherChannel	! interface Port-Channel 1 description CONNECTED TO L2 ACCESS switchport trunk allowed vlan 11 evpn ethernet-segment auto lacp df-election wait-time 1 !	! interface Port-Channel 1 description CONNECTED TO L2 ACCESS switchport trunk allowed vlan 11 evpn ethernet-segment auto lacp df-election wait-time 1 !

Underlay: fabric core and BGP peering

Enterprise campus core networks with a solid underlay network foundation is essential to build highly scalable and resilient BGP EVPN VXLAN fabrics. This section is the second step to build a reliable underlay core network and establish hierarchical iBGP peering among targeted network devices with specific roles.

Note The table is divided into two fabric roles with each step either sharing a common configuration or a unique per-device configuration with a common role.

Step	ES-1 and ES-2	Spine-1 and Spine-2	Border-1 and Border-2
1: Global best practices	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !
2: Underlay interface configuration and best practices	! interface range HundredGig1/0/49-50 description CONNECTED TO SPINE ip ospf 100 area 0 ip ospf network point-to-point carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !	! interface range HundredGig1/0/1-4 description CONNECTED TO CAMPUS CORE NETWORK ip ospf 100 area 0 ip ospf network point-to-point carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out !	! interface range HundredGig1/0/49-50 description CONNECTED TO SPINE ip ospf 100 area 0 ip ospf network point-to-point carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out !

Step	ES-1 and ES-2	Spine-1 and Spine-2	Border-1 and Border-2

3: OSPF routing configuration and best practices	ES-1
	!
	router ospf 100
	 router-id 10.200.255.101
	
	 max-metric router-lsa
	  include-stub
	summary-lsa
	 external-lsa on-startup
	
	  wait-for-bgp
	 nsf cisco
	 fast-reroute per-prefix
	 enable
	  prefix-priority low
	 area 101 stub no-summary
	
	 passive-interface
	default
	 no passive-interface
	Port-Channel 128
	 no passive-interface
	HundredGig1/0/49
	 no passive-interface
	HundredGig1/0/50
	!
	ES-2
	!
	router ospf 100
	 router-id 10.200.255.102
	
	 max-metric router-lsa
	  include-stub
	summary-lsa
	 external-lsa on-startup
	 wait-for-bgp
	 nsf cisco
	 fast-reroute per-prefix
	 enable
	  prefix-priority low
	 area 101 stub no-summary
	
	 passive-interface
	default
	 no passive-interface
	Port-Channel 128
	 no passive-interface
	HundredGig1/0/49
	 no passive-interface
	HundredGig1/0/50
	!
		SPINE-1
		!
		router ospf 100
		 router-id 10.200.255.3
		
		 max-metric router-lsa
		  include-stub
		summary-lsa
		 external-lsa on-startup
		
		  wait-for-bgp
		 nsf cisco
		 fast-reroute per-prefix
		 enable
		  prefix-priority low
		passive-interface default
		
		 no passive-interface
		HundredGig1/0/1
		 no passive-interface
		HundredGig1/0/2
		no passive-interface
		HundredGig1/0/3
		 no passive-interface
		HundredGig1/0/4
		!
		SPINE-2
		!
		router ospf 100
		 router-id 10.200.255.4
		
		 max-metric router-lsa
		  include-stub
		summary-lsa
		 external-lsa on-startup
		
		  wait-for-bgp
		 nsf cisco
		 fast-reroute per-prefix
		 enable
		  prefix-priority low
		 passive-interface
		default
		 no passive-interface
		HundredGig1/0/1
		 no passive-interface
		HundredGig1/0/2
		no passive-interface
		HundredGig1/0/3
		 no passive-interface
		HundredGig1/0/4
		!
			BORDER-1
			!
			router ospf 100
			 router-id 10.200.255.1
			
			 max-metric router-lsa
			  include-stub
			summary-lsa
			 external-lsa on-startup
			
			  wait-for-bgp
			 nsf cisco
			 fast-reroute per-prefix
			 enable
			  prefix-priority low
			passive-interface default
			
			 no passive-interface
			HundredGig1/0/49
			 no passive-interface
			HundredGig1/0/50
			!
			BORDER-2
			!
			router ospf 100
			 router-id 10.200.255.2
			
			 max-metric router-lsa
			  include-stub
			summary-lsa
			 external-lsa on-startup
			
			  wait-for-bgp
			 nsf cisco
			 fast-reroute per-prefix
			 enable
			  prefix-priority low
			
			 passive-interface
			default
			 no passive-interface
			HundredGig1/0/49
			 no passive-interface
			HundredGig1/0/50
			!

Step	ES-1 and ES-2	Spine-1 and Spine-2	Border-1 and Border-2
4: BGP routing configuration	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !
5: Peer- session and peer-policy templates and parameters for leaf switches	! template peer-session EVPN-SPINE-PEER-SESSION-POLICY remote-as 65101 description EVPN-SPINE-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-SPINE-PEER-POLICY send-community both !	! template peer-session EVPN-LEAF-PEER-SESSION-POLICY remote-as 65101 description EVPN-LEAF-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-LEAF-PEER-POLICY route-reflector-client send-community both !	-
6: Peer-session and peer-policy templates and parameters for border switches		! template peer-session EVPN-BORDER-PEER-SESSION-POLICY remote-as 65101 description EVPN-BORDER-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-BORDER-PEER-POLICY route-reflector-client send-community both !	! template peer-session EVPN-SPINE-PEER-SESSION-POLICY remote-as 65101 description EVPN-SPINE-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-SPINE-PEER-POLICY send-community both !
7: Disable intra-cluster EVPN multihome leaf		! no bgp client-to-client reflection intra-cluster cluster-id any !	-

Step	ES-1 and ES-2	Spine-1 and Spine-2	Border-1 and Border-2
8: Configure border-spine iBGP peering		! neighbor 10.200.255.1 inherit peer-session EVPN-BORDER- PEER-SESSION-POLICY ! neighbor 10.200.255.2 inherit peer-session EVPN-BORDER- PEER-SESSION-POLICY !	! neighbor 10.200.255.3 inherit peer-session EVPN-SPINE- PEER-SESSION-POLICY ! neighbor 10.200.255.4 inherit peer-session EVPN-SPINE- PEER-SESSION-POLICY !
9: Configure leaf iBGP peering	! neighbor 10.200.255.3 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY ! neighbor 10.200.255.4 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY !	! neighbor 10.200.255.101 inherit peer-session EVPN-LEAF- PEER-SESSION-POLICY neighbor 10.200.255.101 cluster-id 1.1.1.1 ! neighbor 10.200.255.102 inherit peer-session EVPN-LEAF- PEER-SESSION-POLICY neighbor 10.200.255.102 cluster-id 1.1.1.1 !	

Step	ES-1 and ES-2	Spine-1 and Spine-2	Border-1 and Border-2
10: Activate leaf and border iBGP Peering under L2VPN EVPN AF	! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.3 activate neighbor 10.200.255.3 send-community both neighbor 10.200.255.3 inherit peer-policy EVPN-SPINE-PEER-POLICY neighbor 10.200.255.4 activate neighbor 10.200.255.4 send-community both neighbor 10.200.255.4 inherit peer-policy EVPN-SPINE-PEER-POLICY !	! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.1 activate neighbor 10.200.255.1 send-community both neighbor 10.200.255.1 inherit peer-policy EVPN-BORDER-PEER-POLICY neighbor 10.200.255.2 activate neighbor 10.200.255.2 send-community both neighbor 10.200.255.2 inherit peer-policy EVPN-BORDER-PEER-POLICY ! neighbor 10.200.255.101 activate neighbor 10.200.255.101 send-community both neighbor 10.200.255.101 inherit peer-policy EVPN-LEAF-PEER-POLICY neighbor 10.200.255.102 activate neighbor 10.200.255.102 send-community both neighbor 10.200.255.102 inherit peer-policy EVPN-LEAF-PEER-POLICY !	! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.3 activate neighbor 10.200.255.3 send-community both neighbor 10.200.255.3 inherit peer-policy EVPN-SPINE-PEER-POLICY neighbor 10.200.255.4 activate neighbor 10.200.255.4 send-community both neighbor 10.200.255.4 inherit peer-policy EVPN-SPINE-PEER-POLICY !

Overlay routed networks

The overlay network configuration is the final step to enable the fabric in an enterprise campus. This step-by-step procedure shows VTEPs configured in a routed overlay network that exchanges IP prefixes between external and internal network domains.

On switches, the routed overlay configuration is implemented within EVPN multihoming-based networks. The configuration example also shows the standard procedure to implement and manage the overlay network in other system configuration modes such as routed access, Cisco StackWise Virtual, and First Hop Redundancy Protocol (FHRP).

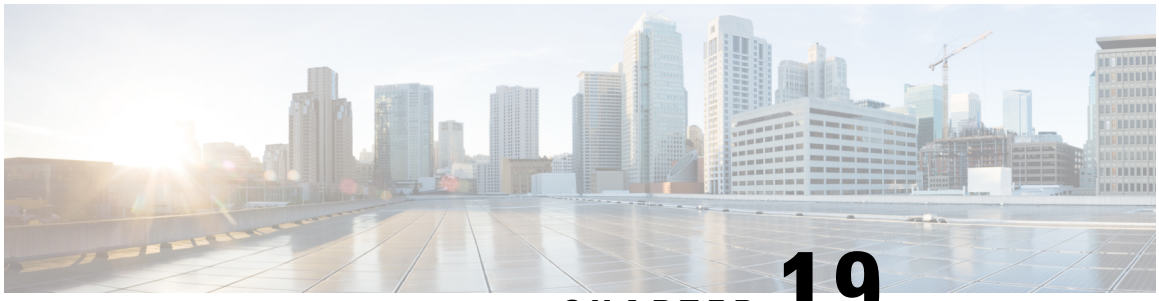
Step	ES-1 and ES-2	Border-1 and Border-2
1: IP-VRF configuration	ES-1 ! vrf definition green rd 10.200.255.101:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-2 ! vrf definition green rd 10.200.255.102:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching !	BORDER-1 ! vrf definition green rd 10.200.255.1:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! BORDER-2 ! vrf definition green rd 10.200.255.2:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching !

2: IP VRF core VLAN	!
	vlan 101
	  name VRF_GREEN_CORE_VLAN
	!
	vlan configuration 101
	 member vni 10011
	!
	interface vlan 101
	 description CORE VLAN – VRF GREEN
	 vrf forwarding green
	 ip unnumbered Loopback0
	 no autostate
	 !
		!
		vlan 101
		  name VRF_GREEN_CORE_VLAN
		
		!
		vlan configuration 101
		 member vni 10011
		!
		interface vlan 101
		 description CORE VLAN –
		VRF GREEN
		 vrf forwarding green
		 ip unnumbered Loopback0
		 no autostate
		 !

3: IP VRF L3VNI to NVE interface binding	! interface nve 1 member vni 10011 vrf green !	! interface nve 1 member vni 10011 vrf green !

Step	ES-1 and ES-2	Border-1 and Border-2
4: Network edge to access or external domain.	! interface Vlan 11 description ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.11.1.254 255.255.255.0 !	BORDER-1 ! interface Vlan 2001 description FIREWALL HANDOFF – VRF GREEN vrf forwarding green ip address 21.1.1.0 255.255.255.254 ! BORDER-2 ! interface Vlan 2002 description FIREWALL HANDOFF – VRF GREEN vrf forwarding green ip address 21.1.1.2 255.255.255.254 !
5: Route-map policy	! route-map SPINE-ROUTE-POLICY-OUT permit 10 description ROUTED OVERLAY NETWORK POLICY match evpn route-type 5 !	
6: Apply spine policy to BGP template	! router bgp 65101 ! template peer-policy EVPN-SPINE-PEER-POLICY route-map SPINE-ROUTE-POLICY-OUT out !	

Step	ES-1 and ES-2	Border-1 and Border-2
7: IP VRF routing	! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected maximum-paths ibgp 2 !	BORDER-1 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.1 remote-as 65001 neighbor 21.1.1.1 activate maximum-paths ibgp 2 ! BORDER-2 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.3 remote-as 65001 neighbor 21.1.1.3 activate maximum-paths ibgp 2 !



CHAPTER 19

Overlay: Distributed Anycast Bridged

The distributed anycast bridged overlay enables efficient Layer 2 connectivity across network segments by using anycast gateway addresses, allowing seamless host mobility and redundancy.

- Overview: Distributed Anycast Gateway bridged overlay networks, on page 727
- Guidelines for DAG bridged overlay networks, on page 728
- Understanding DAG bridged overlay networks, on page 728
- Configure DAG bridged overlay networks, on page 730
- Verify the DAG bridged overlay network configuration, on page 740
- Reference configuration for DAG bridged overlay networks, on page 743
- Layer 2 EVPN multihoming, on page 743
- Underlay: fabric core and BGP peering, on page 749
- Overlay: DAG bridged networks, on page 757

Overview: Distributed Anycast Gateway bridged overlay networks

In modern enterprise campus environments, most application traffic flows between wired and wireless users and between resources that are hosted within on-premises data centers or across cloud. Network administrators build structured IP addressing plans with VLAN and subnet on a per-distribution block and enable segmentation through routed overlay or selectively stretching flood-free IP subnets with Distributed Anycast Gateway Routed VXLAN fabric network.

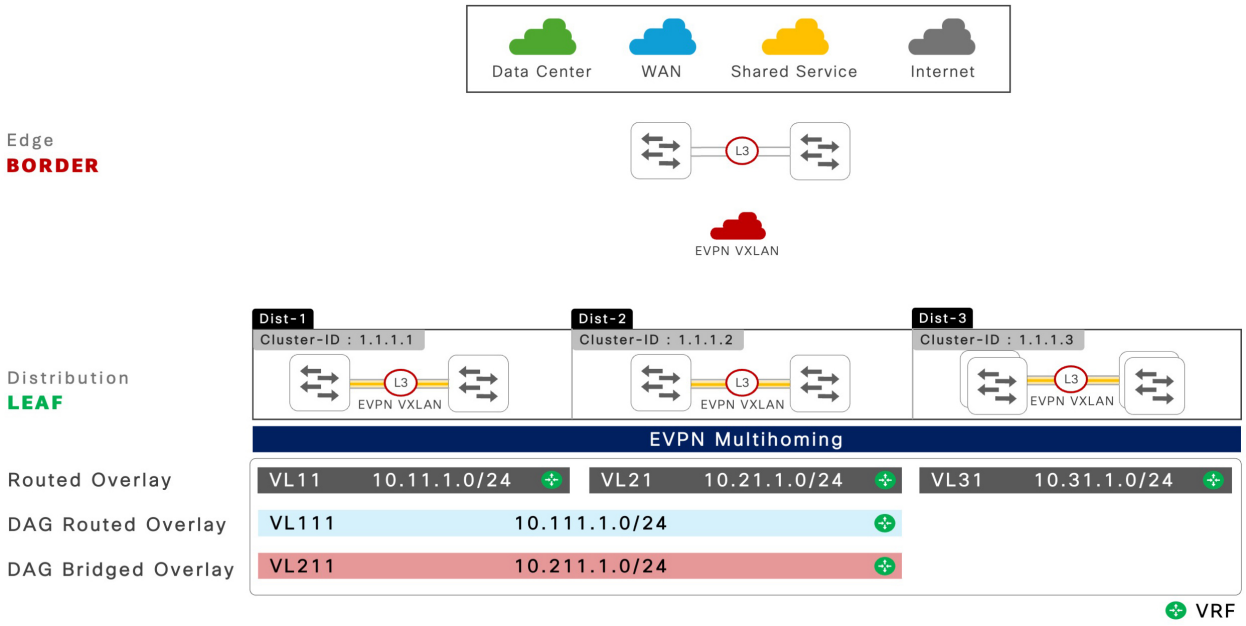
With the digitally evolving IP-everywhere networks, in some scenarios network administrators are mandated to maintain modern fabric networks, while supporting scarce VLAN segments with legacy enterprise network applications and endpoints operating at link-layer without IP addresses. The Distributed Anycast Gateway (DAG) technology supports classic methods to extend the Layer 2 flood-boundary for selective VLANs over the dynamic Layer 2 VXLAN tunnels.

The DAG-bridged overlay networks stretch IP subnets and the Layer 2 VLAN flood-boundary with traditional Integrated Routing and Bridging (IRB) function between targeted switches implemented in EVPN multihoming fabric mode. The DAG-bridged overlay networks extend the Layer 2 flood boundary across fabric networks for all endpoints including hosts with and without IP addresses.

Switches support the coexistence of routed, DAG-routed and DAG-bridged overlay networks that provide the flexibility to scale fabric networks while enabling enterprise application requirements.

The following figure shows a reference EVPN multihoming network with DAG-bridged overlay for extending Layer 2 flood-boundary for a specific VLAN or IP subnet between pairs of switches.

Figure 73: EVPN multihoming with DAG-bridged overlay fabric network



Guidelines for DAG bridged overlay networks

Follow these guidelines when planning for DAG bridged overlay networks.

- Extended bridge-domain boundaries build larger Layer 2 flood domains.
- The DAG bridge mode is suggested for the deployment of selective VLANs or subnets based on technical requirements, such as non-IP and silent hosts.

Refer to the downstream Layer 2 network access device product datasheet to determine the MAC address scale numbers.

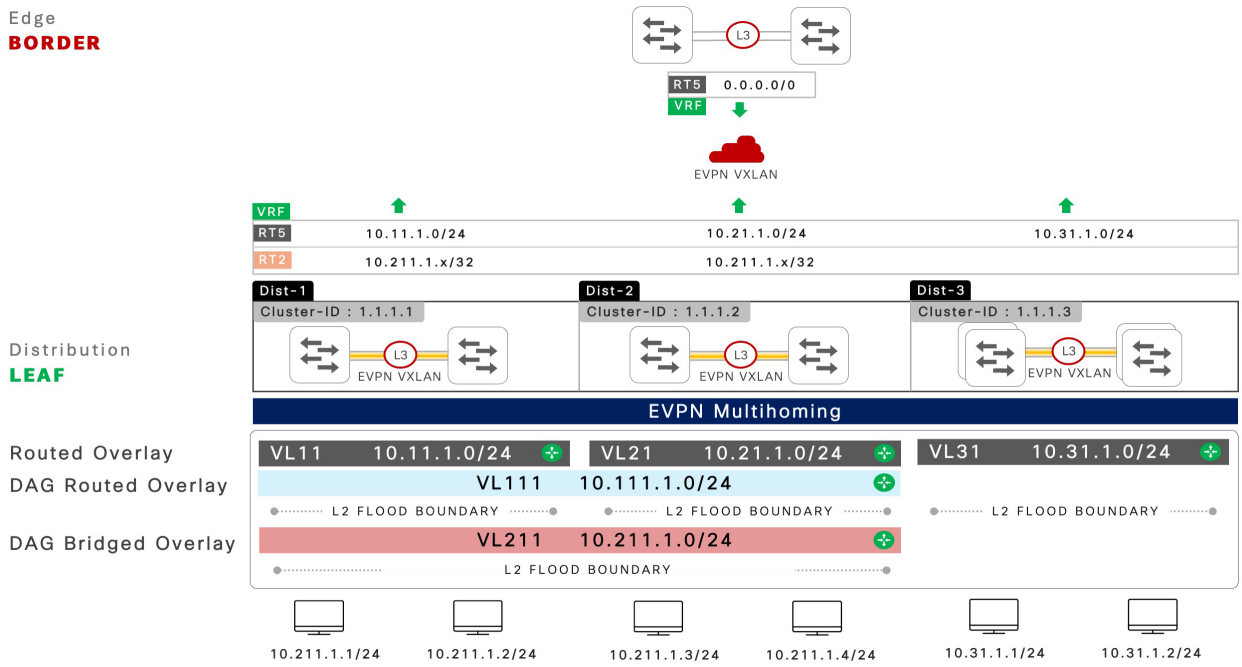
Understanding DAG bridged overlay networks

Based on business requirements, enterprise networks have the flexibility to support unique per-VLAN basis overlay functions. Routed overlay is proven and recommended for the large-scale overlay networks, DAG-routed overlay provides flexible Layer 2 flood-free IP subnet stretch for selective networks, and the DAG bridge overlay addresses non-routable legacy endpoints and applications by extending the Layer 2 network across the fabric.

Switches support the coexistence of routed, Distributed Anycast Gateway-routed, and Distributed Anycast Gateway-bridged overlay networking options within the same macro-segmented IP VRF and the micro-segmented VXLAN GPO-enabled environment.

The following figure illustrates the characteristics of an EVPN multihoming network with DAG bridged overlay fabric.

Figure 74: Key characteristics of a DAG bridged fabric overlay network



Hierarchical BGP control plane

The recommended deployment consideration to support controlled route management between EVPN multihoming peers and spine switches in large-scale campus networks is the two-tier hierarchical BGP control plane. Each BGP peering layer manages domain-specific EVPN routes to enable downstream non-blocking and all-active Layer 2 multipath networks, and manages the advertising of fabric network prefixes to enable secure global network access connectivity in core networks.

EVPN multihoming

Layer 2 networks with EVPN multihoming functions independent of the fabric core EVPN fabric network. A pair of switches in each distribution layer builds and manages EVPN multihoming in each redundancy group.

The unified Layer 2 Ethernet Segment (ES) EtherChannel connection to downstream network device can carry VLANs that are mapped to routed, selective-stretched subnets with DAG-routed overlay and advertised in the fabric core or can continue to be traditional underlay IP networks.

Layer 2 broadcast network boundary

The DAG-bridged overlay networks provide unified Layer 2 or Layer 3 network boundary as the routed overlay and the traditional underlay campus network. The Layer 2 bridge domain and the blast radius remain contained between Layer 2 access and EVPN multihoming-enabled distribution layer systems.

A pair of switches in the distribution layer network dynamically builds the Layer 2 VXLAN tunnel using ingress replication mode to extend Broadcast, Unknown Unicast and Multicast (BUM) traffic. With built-in

loop-detection and prevention techniques, EVPN multihoming-enabled networks support a secure and scalable Layer 2 network over traditional STP based networks.

Stretched IP subnet and gateway redundancy

The DAG routed overlay network enables stretching an IP subnet between targeted EVPN multihoming network devices that support distributed connected IP or non-IP reachability across fabric networks. The selective group of switches in EVPN multihoming mode shares common IPv4 or IPv6 subnet function to support intra-subnet host communication based on the fabric host-routing or MAC address tables while providing load-sharing and gateway redundancy with Anycast gateway as a routed overlay.

Spine route policy

The IPv4 or IPv6 subnets and Layer 2 flood domains stretch across multiple IP gateway systems and demand a non-condensed host-level MAC-only or MAC-IP prefix to support seamless data communication across fabrics and beyond.

Large scale EVPN multihoming-enabled fabric networks benefit from the enforcement of the route policy on leaf switches towards the spine switches.

The iBGP or eBGP prefix advertisement to spine switches in the fabric core is limited to the advertising of RT-2 host prefixes and RT-5 network prefixes along with EAD per-ES RT-1 routes. Switches in EVPN multihoming mode retain unfiltered routing with an iBGP peering session with remote ES switches for non-blocking Layer 2 multipath EtherChannel.

Note Routed overlay networks are recommended for better scale, performance, and resiliency.

Configure DAG bridged overlay networks

This section provides sequential configuration tasks to successfully implement a DAG bridged overlay fabric network.

The BGP EVPN VXLAN-enabled fabric core network must be built with a reliable underlay IP and Layer 2 EVPN multihoming network as prerequisites.

Configure IP VRFs

This task shows how to configure a logically segmented virtual network with an IP VRF to exchange IPv4 or IPv6 network prefixes across all VTEPs in the fabric core network.

Note This configuration task is optional if IP VRFs are already configured on switches as part of the routed overlay network configuration.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition green`	Configures a new virtual routing and forwarding (VRF) instance and enters VRF configuration mode.
Step 4	**rd** *vpn-route-distinguisher* **Example:** `Device(config-vrf)# rd 10.200.255.101:101`	Configures a route distinguisher with a distinct value across all fabric devices.
Step 5	**address-family ipv4** [**multicast** | **unicast**] **Example:** `Device(config-vrf-af)# address-family ipv4 unicast`	Enters IPv4 unicast address-family configuration mode.
Step 6	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **Example:** `Device(config- vrf-af)# route-target 65101:101`	Configures an overlay IPv4 import and export routing policy that matches the extended community value. • The default is **both** and will auto-generate both import and export command line with the configured community value.
Step 7	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **stitching** **Example:** `Device(config- vrf-af)# route-target 65101:101 stitching`	Configures an IPv4 routing import and export overlay policy that matches the extended community value and enables extended NLRI types for BGP EVPN VXLAN-enabled networks. • The default is **both**
Step 8	**address-family ipv6** [**multicast** | **unicast**] **Example:** `Device(config- vrf-af)# address-family ipv6 unicast`	Enters IPv6 unicast address-family configuration mode.

	Command or Action	Purpose
Step 9	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **Example:** `Device(config- vrf-af)# route-target 65101:101`	Configures an IPv6 routing import and export overlay policy that matches the extended community value. • The default is **both** and will auto-generate both import and export command line with the configured community value.
Step 10	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **stitching** **Example:** `Device(config- vrf-af)# route-target 65101:101 stitching`	Configures an IPv6 routing import and export overlay policy that matches the extended community value and enables extended NLRI types for BGP EVPN VXLAN-enabled networks. • The default is **both**.
Step 11	**end** **Example:** `Device(config- vrf-af)# end`	Exits IPv6 unicast address-family configuration mode and returns to privileged EXEC mode.

Configure core VLAN and SVI interfaces

This task shows how to configure core VLAN and SVI interfaces for each of the IP VRFs used for VXLAN encapsulation and decapsulation functions between the downstream non-VXLAN network and remote VXLAN-enabled VTEPs. The task also shows how to configure the NVE interface and bind the L3VNI ID to a logical interface.

Note Skip this task if you have already configured core VLAN and SVI interfaces as part of the routed overlay network configuration.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan** *id* **Example:**	Configures a dedicated VLAN ID for IP VRF core routing. This VLAN must not be used for

	Command or Action	Purpose
	`Device(config)# vlan 101`	any other purpose, and it is recommended to be pruned from Layer 2 trunk ports.
Step 4	**vlan configuration** *id* **Example:** `Device(config)# vlan configuration 101`	Configures a core VLAN to assign the Layer 3 VNI ID and enters VLAN configuration mode.
Step 5	**member vni** *l3vni-id* **Example:** `Device(config-vlan)# member vni 11101`	Configures a unique Layer 3 VNI ID for a specific IP VRF. For example, VRF green has the dedicated VLAN, 101 and Layer 3 VNI, and these IDs must not be shared with other IP VRFs on the same system. • It is recommended that all other VTEPs must have a Layer 3 VNI ID for the same IP VRF.
Step 6	**exit** **Example:** `Device(config-vlan)# exit`	Exits VLAN configuration mode and returns to global configuration mode
Step 7	**interface vlan** *id* **Example:** `Device(config)# interface vlan 101`	Configures a core VLAN SVI interface and assigns it to an IP VRF. • The core VLAN SVI MAC address is used to support routed MAC (RMAC) for Layer 3 overlays.
Step 8	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Maps the IP VRF to a core VLAN SVI interface.
Step 9	**ip unnumbered** *id* **Example:** `Device(config-if)# ip unnumbered Loopback0`	Configures an unnumbered IPv4 address from an underlay logical interface, for example, Loopback0.
Step 10	**no autostate** **Example:** `Device(config-if)# no autostate`	Disables autostate on the interface. • In EVPN deployments, a VLAN used for a core-facing SVI, should not be allowed in any trunks. For a core-facing SVI to function correctly, the **no autostate** command must be configured for the SVI.

	Command or Action	Purpose
Step 11	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.
Step 12	**interface nve** *id* **Example:** `Device(config)# interface nve 1`	Configures a virtual NVE interface to bind one or more L2VNIs and enables system-wide VXLAN forwarding function from the L3VNIs. • Enters interface configuration mode.
Step 13	**source-interface** *id* **Example:** `Device(config-if)# source-interface Loopback0`	Configures a Loopback interface as the source to enable data communication over a VXLAN tunnel.
Step 14	**host-reachability protocol bgp** **Example:** `Device(config-if)# host-reachability protocol bgp`	Configures BGP as the host-reachability control plane protocol on this interface.
Step 15	**member vni** *l3vni-id* **vrf** *vrf-name* **Example:** `Device(config-if)# member vni 11101 vrf green`	Binds the L3VNI to the IP VRF under the NVE interface.
Step 16	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and returns to privileged EXEC mode.

Configure a DAG bridged overlay MAC-VRF

This task shows how to configure a MAC VRF and route-target policy to selectively exchange prefixes between a pair of EVPN multihoming-enabled switches in a single cluster.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:**	Enters global configuration mode.

	Command or Action	Purpose
	`Device# configure terminal`	
Step 3	**l2vpn evpn instance** *id* **vlan-based** **Example:** `Device(config)# l2vpn evpn instance 211 vlan-based`	Configures a L2VPN EVPN virtual instance with a unique ID. Typically, the ID value is the same as the VLAN ID, for example, data VLAN ID 211. • Enters EVPN instance configuration mode.
Step 4	**encapsulation vxlan** **Example:** `Device(config-evpn-evi)# encapsulation vxlan`	Enables VXLAN encapsulation to bridge inter-ES data between two systems.
Step 5	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **Example:** `Device(config-evpn-evi)# route-target both 1.1.1.1:211`	Configures a MAC or IP routing import and export policy that matches the extended community value. • The recommended format for EVPN multihoming cluster VLAN ID is Cluster-1 ID with 1.1.1.1 and DAG-routed VLAN ID 211 can result in the MAC VRF route-target 1.1.1.1:211. The default is **both** and it auto-generates the import and export command lines with the configured community value.
Step 6	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **Example:** `Device(config-evpn-evi)# route-target import 1.1.1.2:211`	Configures inter-cluster DAG bridging with a MAC or IP routing import-only policy that matches the extended community value. • The recommended format for EVPN multihoming cluster VLAN ID is Cluster-1 ID with 1.1.1.2 and DAG-routed VLAN ID 211 can result in the MAC VRF route-target 1.1.1.2:211.
Step 7	**no auto-route-target** **Example:** `Device(config-evpn-evi)# no auto-route-target`	Disables auto route target with the manual route target configured in Step 5 for the DAG-routed overlay MAC VRF. • The default is auto-generated route target in ASN:L2VPN instance ID format. Default enabled with auto-generated RT in ASN:L2VPN Instance ID format.

	Command or Action	Purpose
Step 8	**advertise mac enable** **Example:** `Device(config-evpn-evi)# advertise mac enable`	Enables the originating and advertising of the MAC-only route to enable non-IP data communication.
Step 9	**replication-type** {**ingress** | **static**} **Example:** `Device(config-evpn-evi)# replication-type static`	Enables underlay BUM replication type with multicast for scalable fabric networks.
Step 10	**end** **Example:** `Device(config-evpn-evi)# end`	Exits EVPN instance configuration mode and returns to privileged EXEC mode.

Configure access SVI interfaces

This section provides step-by-step instructions to configure switches in EVPN multihoming mode.

The DAG bridge mode provides a single-unified IP gateway system through multiple leaf network devices. The intra subnet IP or non-IP host communication follows standard remote MAC-to-IP binding resolution, while all inter subnet IP communication is routed through the anycast gateway IP address.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *id* **Example:** `Device(config)# interface Vlan 211`	Configures a Layer 3 SVI interface ID and enters interface configuration mode.
Step 4	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Maps the IP VRF to a core VLAN SVI interface.
Step 5	**ip address** *ip-address mask* **Example:**	Assigns an IPv4 gateway address and mask for selected VLANs.

	Command or Action	Purpose
	`Device(config-if)# ip address 10.211.1.254 255.255.255.0`	• The anycast gateway is auto enabled with common addresses between ES systems.
Step 6	**ipv6 address** *ipv6-address mask* **Example:** `Device(config-if)# ipv6 address 2001:10:211:1::f/64`	Assigns an IPv6 gateway address and prefix for selected VLANs. • The anycast gateway is auto enabled with common addresses between ES systems.
Step 7	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and returns to privileged EXEC mode.

Configure a spine policy for DAG bridged overlay networks

This task provides step-by-step instructions to configure a pair of ES switches with an outbound spine EVPN route policy. The route policy enables controlled EVPN route advertisement by a pair of ES switches, independent of iBGP or eBGP spine-peering session type.

The spine route-map policy to permit RT-5 prefix for routed overlay can be expanded with a new sequence to limit switches advertising the Route-Type 1 and Route-Type 2 (that includes MAC-only and MAC-IP) prefixes for network-wide host route propagation based on DAG MAC VRF route target values.

The iBGP peering with EVPN multihoming system unfiltered EVPN route exchange to build and maintain non-blocking Layer 2 multipath network.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**ip extcommunity-list** [**expanded** *list-name* {**permit** | **deny**} [*regular-expression*] | **standard-lit-number** *standard-list-name* {**permit** | **deny**} [**rt** *extended-community-value*] [**soo** *extended-community-value*] **Example:**	Configures an extended IP community list to permit the DAG-bridge MAC VRF route-target values. • A single IP extended community can be configured with a single RT value or can be configured through regular expressions that match multiples. For example, 21[1-9] for 211-219 DAG-bridged VLAN ID range.

	Command or Action	Purpose
	`Device(config)# ip extcommunity-list expanded DAG-BRIDGED-OVERLAY-EXTCOMM permit 1.1.1.1:211`	
Step 4	**route-map** *name* [**permit** | **deny**] *sequence-number* **Example:** `Device(config)# route-map SPINE-ROUTE-POLICY-OUT permit 30`	Creates a route-map rule with permit or deny match criteria and enters route-map configuration mode. • The optional *sequence-number* argument indicates the order of processing of the route-map rule.
Step 5	**description** *description* **Example:** `Device(config-route-map)# description DAG BRIDGED OVERLAY – VL211`	(Optional) Adds a description for the route-map sequence.
Step 6	**match evpn route-type** [**1-8** | **2-mac-ip** | **2-mac-only**] **Example:** `Device(config-route-map)# match evpn route-type 1`	Configures an EVPN route policy sequence to permit advertising the local EAD per Ethernet Segment (ES) prefixes RT-1) for each local ES.
Step 7	**match evpn route-type** [**1-8** | **2-mac-ip** | **2-mac-only**] **Example:** `Device(config-route-map)# match evpn route-type 2`	Configures an EVPN route policy sequence to permit advertising the local IPv4 or IPv6 host and MAC prefixes (RT-2) mapped to any IP VRF.
Step 8	**exit** **Example:** `Device(config-route-map)# exit`	Exits route-map configuration mode and returns to global configuration mode.
Step 9	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 65101`	Configures BGP autonomous system numbers using 2-bytes or 4-bytes in asplain and asdot formats and enters router configuration mode.
Step 10	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Enters BGP L2VPN address-family configuration mode.
Step 11	**neighbor** *spine-loopback-address* **route-map** *name* **out** **Example:** `Device(config-router-af)# neighbor 10.200.255.3 route-map SPINE-ROUTE-POLICY-OUT out`	Applies outbound route-map policy to each spine peer to limit the advertising of the RT-5 prefix.

	Command or Action	Purpose
Step 12	**end** **Example:** `Device(config-router-af)# end`	Exits BGP L2VPN address-family configuration mode and returns to privileged EXEC mode.

Configure IP VRF address family routing under BGP

This task shows how to advertise network prefixes in IPv4 and IPv6 address-family configurations.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 65101`	Configures BGP autonomous system numbers using 2-bytes or 4-bytes in asplain and asdot formats and enters router configuration mode.
Step 4	**address-family ipv4 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv4 vrf green`	Enables IPv4 address-family for a VRF instance to enable IPv4 network prefix routing in overlay.
Step 5	**advertise l2vpn evpn** **Example:** `Device(config-route-af)# advertise l2vpn evpn`	Enables advertising the IPv4 network prefix routing between IPv4 VRF instances and EVPN fabric core networks.
Step 6	**redistribute connected** [**metric** |**route-map**] *name* **Example:** `Device(config-route-af)# redistribute connected`	Configures IPv4 address-family to advertise all local connected IPv4 network prefixes in EVPN fabric network. • Network prefixes can be advertised by the **network** command as an alternative.
Step 7	**exit** **Example:** `Device(config-route-af)# exit`	Exits IPv4 address-family configuration mode and returns to router configuration mode

	Command or Action	Purpose
Step 8	**address-family ipv6 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv6 vrf green`	Enables IPv6 address-family for a VRF instance to enable IPv6 network prefix routing in overlay.
Step 9	**advertise l2vpn evpn** **Example:** `Device(config-route-af)# advertise l2vpn evpn`	Enables IPv6 network prefix routing between IPv6 VRF instances and EVPN fabric core networks.
Step 10	**redistribute connected** [**metric** |**route-map**] *name* **Example:** `Device(config-route-af)# redistribute connected`	Configures an IPv6 address-family to advertise all local connected IPv6 network prefixes in EVPN fabric network. • Network prefixes can be advertised by the **network** command, as an alternative.
Step 11	**end** **Example:** `Device(config-route-af)# end`	Exits IPv6 address-family configuration mode and returns to privileged EXEC mode.

Verify the DAG bridged overlay network configuration

This section provides examples to verify the DAG-bridged overlay network configuration and states of VTEPs.

The command output may be truncated to focus on critical information for a Day-2 operation and troubleshooting.

IP VRF networks: Verifies the locally configured IP VRF and the associated physical or logical interface bindings to a virtual network on VTEPs including leaf and border devices. The interface binding includes the network edge VLANs connecting to endpoints, and network devices, along with a single core VLAN (VLAN 101) of an IP VRF.

```
ES-1# show vrf green

  Name      Default RD               Protocols     Interfaces
  green     10.200.255.101:101 ipv4     Vl111
                                                   Vl211
                                                   Vl101
BORDER-1# show vrf green

  Name     Default RD              Protocols Interfaces
  green    10.200.255.1:101  ipv4       Vl101
                                                   Vl1101
```

BGP L2VPN EVPN neighbors: Verifies that the two-tier hierarchical iBGP sessions between switches in EVPN multihoming mode, and iBGP peering to a pair of spine systems are operational.

The following command output shows iBGP peering with a pair of spines switches, 10.100.255.3 and 10.100.255.4 and direct iBGP peering between ES-1, local 10.100.255.101 and ES-2, 10.100.255.102 switches in EVPN multihoming mode is operational.

```
ES-1# show bgp l2vpn evpn all summary

BGP router identifier 10.100.255.101, local AS number 65101
<snip>

Neighbor        V    AS     MsgRcvd  MsgSent   TblVer  InQ   OutQ    Up/Down    State/PfxRcd

10.100.255.3     4    65101    18        20      104     0      0     00:04:35         2
10.100.255.4     4    65101    23        25      106     0      0     00:05:19         2
10.100.255.102   4    65101    51        65      104     0      0     00:04:26         28
```

Spine policy: Verifies the VTEP in leaf or border roles implemented in EVPN multihoming mode is configured with a route-map and an outbound policy is applied to each spine, 10.200.255.3.

The spine policy configuration on border switches is optional and not required if the switches are not implemented in EVPN multihoming mode to connect to external network devices, such as firewalls.

```
ES-1# show bgp l2vpn evpn neighbor 10.200.255.3 policy

Neighbor: 10.200.255.3, Address-Family: L2VPN E-VPN
 Locally configured policies:
  send-community extended
 Inherited polices:
  route-map EVPN-SPINE-ROUTE-POLICY-OUT out
```

L2VPN advertised routes: Verifies the VTEP in leaf or border roles implemented in EVPN multihoming mode is advertising the EVPN prefixes to the spine, 10.200.255.3, based on the applied route-map policy.

The following command output from leaf and border switches confirms the advertising of RT-2-MAC-IP EVPN prefixes, and that no additional EVPN route-types are advertised to the spine switches.

```
ES-1# show bgp l2vpn evpn neighbor 10.200.255.3 advertised-route

BGP table version is 7179, local router ID is 10.200.255.101
<snip>
     Network          Next Hop            Metric  LocPrf  Weight  Path
 Route Distinguisher: 10.200.255.101:111
 *>   [2][10.200.255.101:111][0][48][00CCFC4ED1C3][32][10.111.1.1]/24
                       0.0.0.0                  32768     ?
Route Distinguisher: 10.200.255.101:111
 *>   [2][10.200.255.101:111][0][48][00CCFC4ED1C4][32][10.111.1.2]/24
                       0.0.0.0                        32768     ?


Border-1# show bgp l2vpn evpn neighbors 10.200.255.3 advertised-routes

BGP table version is 2895, local router ID is 10.200.255.1
<snip>
     Network          Next Hop            Metric    LocPrf     Weight   Path
Route Distinguisher: 10.100.255.1:101 (default for vrf green)
 *>   [5][10.100.255.1:101][0][0][0.0.0.0]/17
                       21.1.1.1              0          0      65001      ?
 *>   [5][10.100.255.1:101][0][16][111.1.0.0]/17
```

```
                    21.1.1.1              0          0       65001       ?
```

MAC-VRF manual route-target: Verifies that the stretched IP-subnet VLAN is limited to import and export of the MAC or IP route matching the local cluster EVPN multihoming peer switch. With auto route-target disabled, the import and export route-target only lists the manual configuration settings.

The following command output confirms that the bidirectional MAC route policy with route-target 1.1.1.1:12 and the default auto-generated ASN instance ID 65101:211 is not listed along with the manual configuration.

```
ES-1# show l2vpn evpn evi 211 detail

EVPN instance:           211 (VLAN Based)
  RD:                    10.200.255.101:211 (auto)
  Import-RTs:            1.1.1.1:211   1.1.1.2:211
  Export-RTs:            1.1.1.1:211
  Per-EVI Label:         none
  State:                 Established
  Replication Type:      Static
 <snip>
```

VRF IP routing table: Verifies that the leaf implemented in EVPN multihoming mode is advertising the EVPN prefixes to the spine switch 10.200.255.3 based on the applied route-map policy.

The following command output from the border switch confirms that the pair of EVPN multihoming leaf systems configured in DAG routed overlay, each with two hosts are advertising RT2-MAC-IP. Border switches have a dual-path ECMP to each host to support load-sharing and redundancy for overlay network prefixes. And leaf switches receive the RT-5 network prefix from the IP routing table of IP VRFs in the border systems that confirm the advertising of RT-5 EVPN prefixes.

```
ES-1# show ip route vrf green bgp

Routing Table: green
<snip>

Gateway of last resort is 10.200.255.2 to network 0.0.0.0
B*    0.0.0.0/0  [200/0] via 10.200.255.1, 00:00:32, Vlan101
                         [200/0] via 10.200.255.2, 00:00:32, Vlan101


Border-1# show ip route vrf green bgp

Routing Table: green
<snip>
Gateway of last resort is 21.1.1.1 to network 0.0.0.0

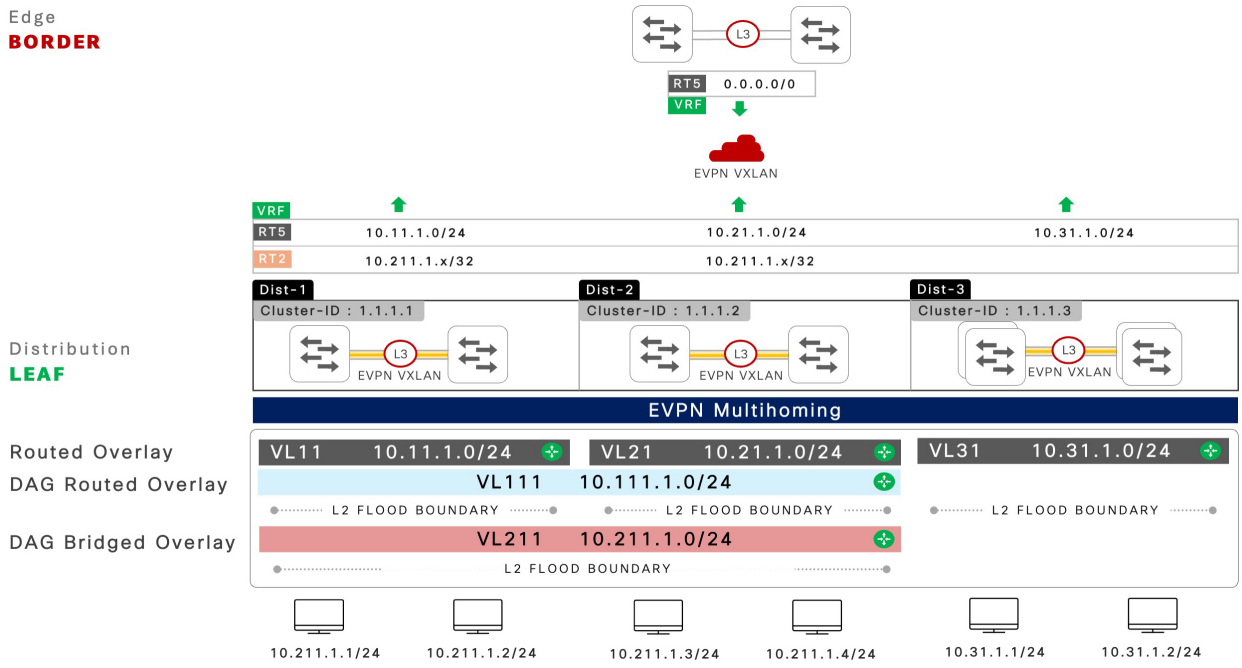
B*    0.0.0.0/0 [20/0] via 21.1.1.1, 6w0d
      10.0.0.0/24 is subnetted, 95 subnets
B        10.111.1.1/32 [200/0] via 10.200.255.102, 00:01:33, Vlan1101
                       [200/0] via 10.200.255.101, 00:01:33, Vlan1101
B        10.111.1.2/32 [200/0] via 10.200.255.102, 00:01:33, Vlan1101
                       [200/0] via 10.200.255.101, 00:01:33, Vlan1101
B        10.111.1.3/32       [200/0] via 10.200.255.104, 00:01:33, Vlan1101
                       [200/0] via 10.200.255.103, 00:01:33, Vlan1101
B        10.111.1.4/32       [200/0] via 10.200.255.104, 00:01:33, Vlan1101
                       [200/0] via 10.200.255.103, 00:01:33, Vlan1101
```

Reference configuration for DAG bridged overlay networks

This section provides a reference network design and configuration to implement a DAG bridged overlay fabric network.

This following figure illustrates a reference network design to support scalable EVPN multihoming and fabric core with hierarchical BGP peering, control-plane peering, and selective stretching of an IP subnet between a pair of switches deployed in the distribution block with overlay fabric network.

Figure 75: EVPN multihoming: DAG-bridged overlay fabric network



This reference configuration shows a three-step process, from the initial network setup to the successful fabric deployment. Network administrators can follow these steps sequentially for initial deployment and to increment the overlay network configuration as the network demands increase.

Layer 2 EVPN multihoming

The base configuration to build a fabric begins with Layer 2 campus networks using EVPN multihoming technology. This section provides the step-by-step configuration on a pair of switches to successfully build a non-blocking all-active Layer 2 network with EVPN multihoming.

Step	ES-1 and ES-2	ES-3 and ES-4
1: Inter-ES Layer 3 EtherChannel	ES-1 ! interface Port-Channel 128 description CONNECTED TO ES PEER no switchport ip address 10.0.0.0 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 101 ip ospf 100 cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking ! ES-2 ! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-2 no switchport ip address 10.0.0.1 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 101 ip ospf 100 cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !	ES-3 ! interface Port-Channel 128 description CONNECTED TO ES PEER no switchport ip address 10.0.0.2 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 102 ip ospf 100 cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking ! ES-4 ! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-1 no switchport ip address 10.0.0.3 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 102 ip ospf 100 cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !

Step	ES-1 and ES-2	ES-3 and ES-4

2: IGP routing and core interface	ES-1
	!
	router ospf 100
	 router-id 10.200.255.101
	 max-metric router-lsa
	  include-stub summary-lsa
	  external-lsa on-startup wait-for-bgp
	
	 nsf cisco
	 fast-reroute per-prefix enable
	  prefix-priority low
	 area 101 stub no-summary
	 passive-interface default
	 no passive-interface Port-Channel 128
	
	 no passive-interface HundredGig1/0/49
	
	 no passive-interface HundredGig1/0/50
	
	!
	interface Loopback 0
	 ip address 10.100.255.101
	  255.255.255.255
	 ip ospf 100 area 0
	!
	ES-2
	!
	router ospf 100
	 router-id 10.200.255.102
	 max-metric router-lsa include-stub
	  summary-lsa
	  external-lsa on-startup wait-for-bgp
	
	 nsf cisco
	 fast-reroute per-prefix enable
	  prefix-priority low
	 area 101 stub no-summary
	 passive-interface default
	 no passive-interface Port-Channel 128
	
	 no passive-interface HundredGig1/0/49
	
	 no passive-interface HundredGig1/0/50
	
	!
	interface Loopback 0
	 ip address 10.100.255.102
	255.255.255.255
	 ip ospf 100 area 0
	!
		ES-3
		!
		router ospf 100
		 router-id 10.200.255.103
		 max-metric router-lsa include-stub
		  summary-lsa
		  external-lsa on-startup
		wait-for-bgp
		 nsf cisco
		 fast-reroute per-prefix enable
		  prefix-priority low
		 area 102 stub no-summary
		 passive-interface default
		 no passive-interface Port-Channel
		128
		 no passive-interface
		HundredGig1/0/49
		 no passive-interface
		HundredGig1/0/50
		!
		interface Loopback 0
		 ip address 10.100.255.103
		255.255.255.255
		 ip ospf 100 area 0
		!
		ES-4
		!
		router ospf 100
		 router-id 10.200.255.104
		 max-metric router-lsa include-stub
		  summary-lsa
		  external-lsa on-startup
		wait-for-bgp
		 nsf cisco
		 fast-reroute per-prefix enable
		  prefix-priority low
		 area 102 stub no-summary
		 passive-interface default
		 no passive-interface Port-Channel
		128
		 no passive-interface
		HundredGig1/0/49
		 no passive-interface
		HundredGig1/0/50
		!
		interface Loopback 0
		 ip address 10.100.255.104
		255.255.255.255
		 ip ospf 100 area 0
		!

Step	ES-1 and ES-2	ES-3 and ES-4
3: iBGP routing	ES-1 ! router bgp 65101 template peer-policy ES-PEER-POLICY send-community both ! template peer-session ES-PEER-SESSION-POLICY remote-as 65101 description EVPN-MH-DIST-2-PEER update-source Loopback0 fall-over host-route ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast ES-PEER-SESSION-POLICY ! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.100.255.102 activate neighbor 10.100.255.102 send-community both neighbor 10.100.255.102 inherit peer-policy ES-PEER-POLICY ! ES-2 ! router bgp 65101 template peer-policy ES-PEER-POLICY send-community both ! template peer-session ES-PEER-SESSION-POLICY remote-as 65101 description EVPN-MH-DIST-2-PEER update-source Loopback0 fall-over host-route ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast neighbor 10.100.255.101 inherit peer-session ES-PEER-SESSION-POLICY ! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.100.255.101 activate neighbor 10.100.255.101 send-community both neighbor 10.100.255.101 inherit peer-policy ES-PEER-POLICY !	

Step	ES-1 and ES-2	ES-3 and ES-4
		ES-3 ! router bgp 65101 template peer-policy ES-PEER-POLICY send-community both ! template peer-session ES-PEER-SESSION-POLICY remote-as 65101 description EVPN-MH-DIST-1-PEER update-source Loopback0 fall-over host-route ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast neighbor 10.100.255.104 inherit peer-session ES-PEER-SESSION-POLICY ! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.100.255.104 activate neighbor 10.100.255.104 send-community both neighbor 10.100.255.104 inherit peer-policy ES-PEER-POLICY ! ES-4 ! router bgp 65101 template peer-policy ES-PEER-POLICY send-community both ! template peer-session ES-PEER-SESSION-POLICY remote-as 65101 description EVPN-MH-DIST-1-PEER update-source Loopback0 fall-over host-route ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast neighbor 10.100.255.103 inherit peer-session ES-PEER-SESSION-POLICY ! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.100.255.103 activate neighbor 10.100.255.103 send-community both neighbor 10.100.255.103 inherit peer-policy

Step	ES-1 and ES-2	ES-3 and ES-4
		ES-PEER-POLICY !
4: Global L2VPN	! l2vpn evpn advertise mac disable anycast-gateway mac auto multicast advertise sync-only multihoming aliasing disable multihoming peering adjacent replication-type ingress router-id Loopback 0 !	! l2vpn evpn advertise mac disable anycast-gateway mac auto multicast advertise sync-only multihoming aliasing disable multihoming peering adjacent replication-type ingress router-id Loopback 0 !
5: Routed VLAN and MAC VRF	! vlan 11 name ROUTED_DATA_VLAN ! l2vpn evpn instance 11 vlan-based encapsulation vxlan ! vlan configuration 11 member evpn-instance 11 vni 11011 ! interface nve 1 source-interface Loopback 0 host-reachability protocol bgp member vni 11011 ingress-replication !	! vlan 11 name ROUTED_DATA_VLAN ! l2vpn evpn instance 11 vlan-based encapsulation vxlan ! vlan configuration 11 member evpn-instance 11 vni 11011 ! interface nve 1 source-interface Loopback 0 host-reachability protocol bgp member vni 11011 ingress-replication !
6: DAG-bridged VLAN and MAC VRF	! vlan 211 name DAG_BRIDGED_DATA_VLAN ! l2vpn evpn instance 211 vlan-based encapsulation vxlan route-target 1.1.1.1:211 route-target import 1.1.1.2:211 no auto-route-target replication-type static advertise mac enable ! vlan configuration 12 member evpn-instance 12 vni 11211 ! interface nve 1 member vni 11211 mcast-group 239.1.1.1 !	! vlan 211 name DAG_BRIDGED_DATA_VLAN ! l2vpn evpn instance 211 vlan-based encapsulation vxlan route-target 1.1.1.2:211 route-target import 1.1.1.1:211 no auto-route-target replication-type static advertise mac enable ! vlan configuration 12 member evpn-instance 12 vni 11211 ! interface nve 1 member vni 11211 mcast-group 239.1.1.1 !

Step	ES-1 and ES-2	ES-3 and ES-4
7: ES EtherChannel	! interface Port-Channel 1 description CONNECTED TO L2 ACCESS switchport trunk allowed vlan 11,211 evpn ethernet-segment auto lacp df-election wait-time 1 !	! interface Port-Channel 1 description CONNECTED TO L2 ACCESS switchport trunk allowed vlan 11,211 evpn ethernet-segment auto lacp df-election wait-time 1 !

Underlay: fabric core and BGP peering

Enterprise campus core networks with solid underlay network foundation is the key for highly scalable, resilient BGP EVPN VXLAN fabric networks. This section is the second step to build a reliable underlay core network for fabric and hierarchical BGP peering on targeted network devices with specific roles.

Note The table is subdivided into two fabric roles with each step either sharing common configuration or a unique per-device with a common role.

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
1: Global best practices	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
2: Underlay interface configuration and best practices	! interface range HundredGig1/0/49-50 description CONNECTED TO SPINE ip ospf 100 area 0 ip ospf network point-to-point ip pim sparse-mode carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !	! interface range HundredGig1/0/1-4 description CONNECTED TO CAMPUS CORE NETWORK ip ospf 100 area 0 ip ospf network point-to-point ip pim sparse-mode carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out !	! interface range HundredGig1/0/49-50 description CONNECTED TO SPINE ip ospf 100 area 0 ip ospf network point-to-point ip pim sparse-mode carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
3: OSPF routing configuration and best practices	ES-1 and ES-2 ! router ospf 100 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low area 101 stub no-summary passive-interface default no passive-interface Port-Channel 128 no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 ! ES-3 and ES-4 ! router ospf 100 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low area 102 stub no-summary passive-interface default no passive-interface Port-Channel 128 no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 !	SPINE-1 ! router ospf 100 router-id 10.200.255.3 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/1 no passive-interface HundredGig1/0/2 no passive-interface HundredGig1/0/3 no passive-interface HundredGig1/0/4 ! SPINE-2 ! router ospf 100 router-id 10.200.255.4 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/1 no passive-interface HundredGig1/0/2 no passive-interface HundredGig1/0/3 no passive-interface HundredGig1/0/4 !	BORDER-1 ! router ospf 100 router-id 10.200.255.1 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 ! BORDER-2 ! router ospf 100 router-id 10.200.255.2 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
4: BGP routing configuration and best practices	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !
5: Peer-session and peer-policy templates and parameters for leaf switches	! template peer-session EVPN-SPINE-PEER-SESSION-POLICY remote-as 65101 description EVPN-SPINE-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-SPINE-PEER-POLICY send-community both !	! template peer-session EVPN-LEAF-PEER-SESSION-POLICY remote-as 65101 description EVPN-LEAF-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-LEAF-PEER-POLICY route-reflector-client send-community both !	
6: Peer-session and policy templates and parameters for border switches		! template peer-session EVPN-BORDER-PEER-SESSION-POLICY remote-as 65101 description EVPN-BORDER-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-BORDER-PEER-POLICY route-reflector-client send-community both !	! template peer-session EVPN-SPINE-PEER-SESSION-POLICY remote-as 65101 description EVPN-SPINE-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-SPINE-PEER-POLICY send-community both !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
7: Disable intra-cluster EVPN multihome leaf reflection		! no bgp client-to-client reflection intra-cluster cluster-id any ! !	
8: Border-spine iBGP peering		! neighbor 10.200.255.1 inherit peer-session EVPN-BORDER-PEER-SESSION-POLICY ! neighbor 10.200.255.2 inherit peer-session EVPN-BORDER-PEER-SESSION-POLICY !	! neighbor 10.200.255.3 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY ! neighbor 10.200.255.4 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
9: Leaf iBGP peering	! neighbor 10.200.255.3 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY ! neighbor 10.200.255.4 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY !	! neighbor 10.200.255.101 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.101 cluster-id 1.1.1.1 ! neighbor 10.200.255.102 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.102 cluster-id 1.1.1.1 ! neighbor 10.200.255.103 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.101 cluster-id 1.1.1.2 ! neighbor 10.200.255.104 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.102 cluster-id 1.1.1.2 !	

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
10: Activate leaf and border iBGP peering under L2VPN EVPN address family	! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.3 activate neighbor 10.200.255.3 send-community both neighbor 10.200.255.3 inherit peer-policy EVPN-SPINE-PEER-POLICY neighbor 10.200.255.4 activate neighbor 10.200.255.4 send-community both neighbor 10.200.255.4 inherit peer-policy EVPN-SPINE-PEER-POLICY !		! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.3 activate neighbor 10.200.255.3 send-community both neighbor 10.200.255.3 inherit peer-policy EVPN-SPINE-PEER-POLICY neighbor 10.200.255.4 activate neighbor 10.200.255.4 send-community both neighbor 10.200.255.4 inherit peer-policy EVPN-SPINE-PEER-POLICY !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
		! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.1 activate neighbor 10.200.255.1 send-community both neighbor 10.200.255.1 inherit peer-policy EVPN-BORDER-PEER-POLICY neighbor 10.200.255.2 activate neighbor 10.200.255.2 send-community both neighbor 10.200.255.2 inherit peer-policy EVPN-BORDER-PEER-POLICY ! neighbor 10.200.255.101 activate neighbor 10.200.255.101 send-community both neighbor 10.200.255.101 inherit peer-policy EVPN-LEAF-PEER-POLICY neighbor 10.200.255.102 activate neighbor 10.200.255.102 send-community both neighbor 10.200.255.102 inherit peer-policy EVPN-LEAF-PEER-POLICY ! neighbor 10.200.255.103 activate neighbor 10.200.255.103 send-community both neighbor 10.200.255.103	

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
		inherit peer-policy EVPN-LEAF-PEER-POLICY neighbor 10.200.255.104 activate neighbor 10.200.255.104 send-community both neighbor 10.200.255.104 inherit peer-policy EVPN-LEAF-PEER-POLICY !	

Overlay: DAG bridged networks

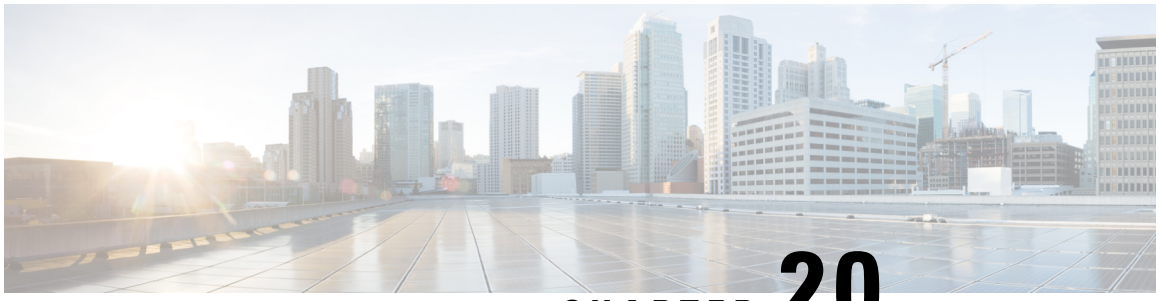
The overlay network configuration is the final step to enable a fabric in the enterprise campus. This section provides step-by-step configuration procedures to be implemented on VTEPs involved in DAG bridged overlay network that exchanges IP prefixes between external and internal network domains.

Note The table is divided between two fabric roles with each step either sharing a common configuration or a unique per-device configuration with a common role.

Step	ES-1, ES-2, ES-3 and ES-4	Border-1 and Border-2
1: IP VRF configuration	ES-1 ! vrf definition green rd 10.200.255.101:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-2 ! vrf definition green rd 10.200.255.102:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-3 ! vrf definition green rd 10.200.255.103:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-4 ! vrf definition green rd 10.200.255.104:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching !	BORDER-1 ! vrf definition green rd 10.200.255.1:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! BORDER-2 ! vrf definition green rd 10.200.255.2:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching !
2: IP VRF core VLAN configuration	! vlan 101 name VRF_GREEN_CORE_VLAN ! vlan configuration 101 member vni 10011 ! interface vlan 101 description CORE VLAN – VRF GREEN vrf forwarding green ip unnumbered Loopback0 no autostate !	! vlan 101 name VRF_GREEN_CORE_VLAN ! vlan configuration 101 member vni 10011 ! interface vlan 101 description CORE VLAN – VRF GREEN vrf forwarding green ip unnumbered Loopback0 no autostate !
3: IP VRF L3VNI to NVE interface binding	! interface nve 1 member vni 10011 vrf green !	! interface nve 1 member vni 10011 vrf green !

Step	ES-1, ES-2, ES-3 and ES-4	Border-1 and Border-2
4: Network edge to access or external domain.	ES-1 and ES-2 ! interface Vlan 11 description ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.11.1.254 255.255.255.0 ! interface Vlan 211 description DAG BRIDGED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.211.1.254 255.255.255.0 ! ES-3 and ES-4 ! interface Vlan 21 description ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.21.1.254 255.255.255.0 ! interface Vlan 211 description DAG BRIDGED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.211.1.254 255.255.255.0 !	BORDER-1 ! interface Vlan 2001 description FIREWALL HANDOFF – VRF GREEN vrf forwarding green ip address 21.1.1.0 255.255.255.254 ! BORDER-2 ! interface Vlan 2002 description FIREWALL HANDOFF – VRF GREEN vrf forwarding green ip address 21.1.1.2 255.255.255.254 !
5: IP extended community matching DAG-bridge MAC VRF route target	ES-1 and ES-2 ! ip extcommunity-list expanded DAG-BRIDGED-OVERLAY-EXTCOMM permit 1.1.1.1:211 ! ES-3 and ES-4 ! ip extcommunity-list expanded DAG-BRIDGED-OVERLAY-EXTCOMM permit 1.1.1.2:211 !	
6: Route map policy	! route-map SPINE-ROUTE-POLICY-OUT permit 10 description ROUTED OVERLAY NETWORK POLICY match evpn route-type 5 ! route-map SPINE-ROUTE-POLICY-OUT permit 30 description DAG BRIDGED OVERLAY NETWORK POLICY match extcommunity DAG-BRIDGED-OVERLAY-EXTCOMM match evpn route-type 1 match evpn route-type 2 !	

Step	ES-1, ES-2, ES-3 and ES-4	Border-1 and Border-2
7: Apply spine policy to BGP template	! router bgp 65101 ! template peer-policy EVPN-SPINE-PEER-POLICY route-map SPINE-ROUTE-POLICY-OUT out !	BORDER-1 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.1 remote-as 65001 neighbor 21.1.1.1 activate maximum-paths ibgp 2 ! BORDER-2 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.3 remote-as 65001 neighbor 21.1.1.3 activate maximum-paths ibgp 2 !
8: IP VRF routing	! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected maximum-paths ibgp 2 !	BORDER-1 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.1 remote-as 65001 neighbor 21.1.1.1 activate maximum-paths ibgp 2 ! BORDER-2 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.3 remote-as 65001 neighbor 21.1.1.3 activate maximum-paths ibgp 2 !



CHAPTER 20

Overlay:Distributed Anycast Gateway Routed

This feature enables distributed anycast gateway routing in overlay networks, allowing efficient traffic forwarding and gateway redundancy.

- Distributed Anycast Gateway: routed overlay network overview, on page 761
- Restrictions for DAG routed overlay networks, on page 762
- Understanding DAG routed overlay networks, on page 762
- Configure DAG routed overlay networks, on page 764
- Verify the DAG routed overlay configuration, on page 773
- Reference configuration for DAG routed overlay networks, on page 776

Distributed Anycast Gateway: routed overlay network overview

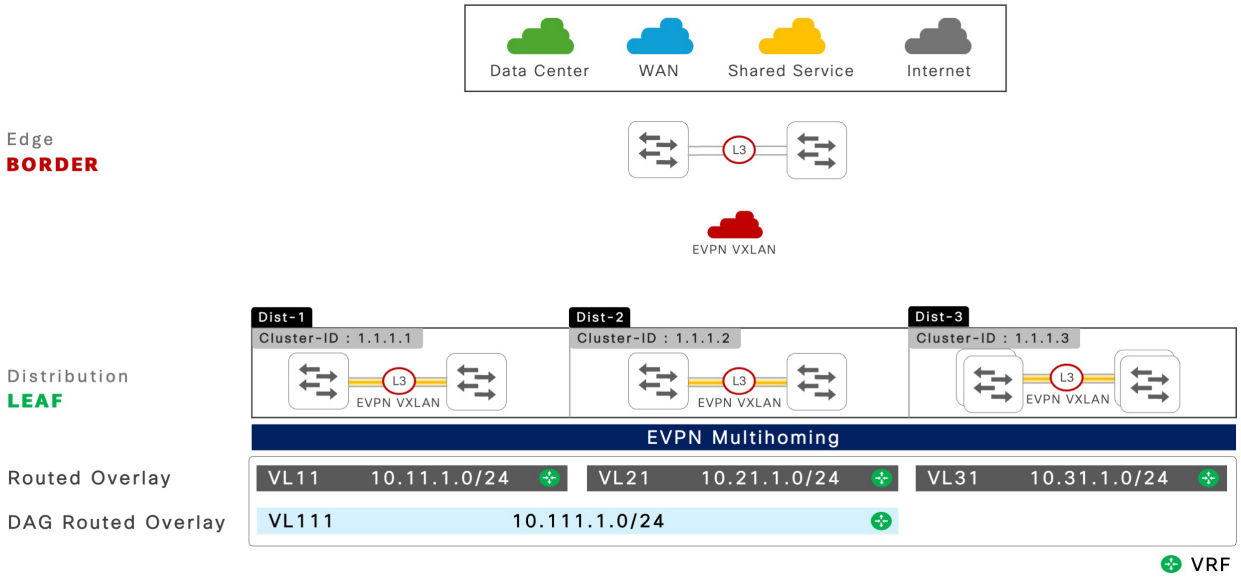
In modern enterprise campus environments, application traffic typically flows between wired and wireless users, and resources hosted either on-premises or across cloud environments. To manage this efficiently, network administrators implement a structured IP addressing plan with VLANs and subnets organized per distribution block. This approach enables segmentation through a routed overlay VXLAN fabric network, which provides scalable, flexible, and secure network segmentation.

In certain cases, the VLAN or subnet networks may require stretching beyond a single IP gateway boundary to support latency-sensitive or legacy enterprise applications. The Distributed Anycast Gateway (DAG) technology enables this by allowing secure and segmented communication between IP endpoints over a dynamic VXLAN tunnel.

With a Layer 2 flood-free network, the DAG-routed overlay network enables the stretching of IP subnets across switches implemented in EVPN multihoming fabric mode across targeted leaf switches. This approach supports scalable and secure network segmentation tailored to enterprise application requirements. Switches support the co-existence of routed and DAG routed overlay networks.

The following illustration shows a deployment scenario for an EVPN multihoming network with the co-existence of routed and DAG-routed overlay networks. The illustration also shows the stretching of a specific IP subnet, VLAN-111 between two targeted switches, one in network distribution Block 1 and the other in Block 2.

Figure 76: EVPN multihoming with DAG-routed overlay fabric network



Restrictions for DAG routed overlay networks

The following restrictions apply to the DAG Routed Overlay feature.

- Non-IP endpoint communication beyond local EVPN multihoming network devices is not supported.
- Silent-host endpoint communication beyond local EVPN multihoming network devices is not supported.

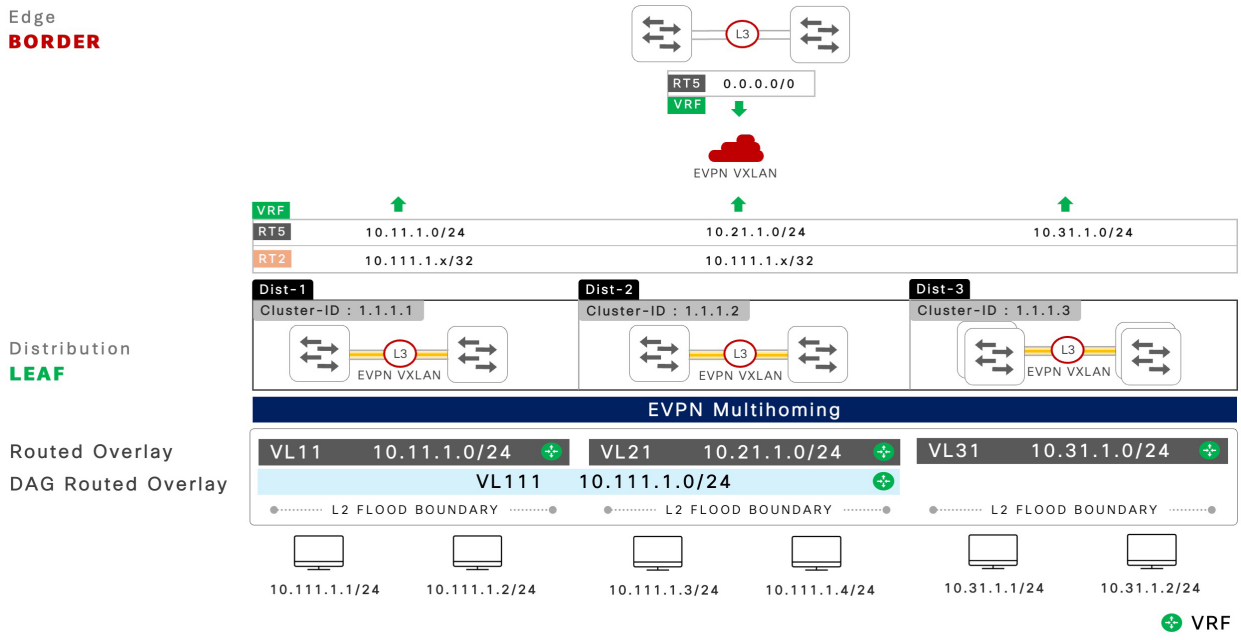
Understanding DAG routed overlay networks

Enterprise networks require flexibility to support unique per-VLAN overlay functions that are tailored to business requirements. Routed overlay is a proven and recommended approach for large-scale overlay networks, which can implement DAG-routed overlay to selectively stretch IP subnets without extending the Layer 2 flood boundary. Network administrators can implement the DAG-routed overlay to stretch an IP subnet across targeted leaf switches without extending the Layer 2 flood boundary.

Switches support the multiple flexible co-existence of routed and distributed Anycast gateway routed overlay networking options within the same macro-segmented IP VRF, and the micro-segmented VXLAN-GPO enabled environment.

The following figure displays the characteristics of an EVPN multihoming network with a DAG-routed overlay fabric.

Figure 77: Key network characteristics of a DAG-routed overlay fabric



Hierarchical BGP control plane

The recommended deployment consideration to support controlled route management between EVPN multihoming peers and spine switches in large-scale campus networks is the two-tier hierarchical BGP control plane. Each BGP peering layer manages domain-specific EVPN routes to enable downstream non-blocking and all-active Layer 2 multipath networks, and manages the advertising of fabric network prefixes to enable secure global network access connectivity in core networks.

EVPN multihoming

Layer 2 networks with EVPN multihoming functions independent of the fabric core EVPN fabric network. A pair of switches in each distribution layer builds and manages EVPN multihoming in each redundancy group.

The unified Layer 2 Ethernet Segment (ES) EtherChannel connection to downstream network device can carry VLANs that are mapped to routed, selective-stretched subnets with DAG-routed overlay and advertised in the fabric core or can continue to be traditional underlay IP networks.

Layer 2 broadcast network boundary

DAG routed overlay networks provide unified Layer 2 or Layer 3 network boundary as the routed overlay and the traditional underlay campus network. The Layer 2 bridge domain and the blast-radius remain contained between Layer 2 access and EVPN multihoming-enabled distribution layer systems.

A pair of switches in the distribution layer network dynamically builds the Layer 2 VXLAN tunnel using ingress replication mode to extend Broadcast, Unknown Unicast and Multicast (BUM) traffic. With built-in loop-detection and prevention techniques, EVPN multihoming-enabled networks support a secure and scalable Layer 2 network over traditional STP based networks.

IPv4 ARP and IPv6 proxy

Switches support IPv4 ARP and IPv6 Neighbor Discovery (ND) and Duplicate Address Detection (DAD) proxy functions on a per VLAN or subnet basis.

The intended, stretched IP subnet interface is configured with local proxy functions on anycast gateway switches. The incoming IPv4 ARP request and IPv6 ND/DAD messages are intercepted, processed, and responded to by the anycast gateway switch with the local anycast MAC address (00:00:5e:00:01:01). As a result, the IP subnet with IPv4 or IPv6 endpoints can be selectively stretched between EVPN multihoming network devices without extending the Layer 2 flood boundary that can impact scale, performance, and security in large scale environments.

Structured IP subnet with proxy and gateway redundancy

The DAG routed overlay network enables stretching an IP subnet between targeted EVPN multihoming network devices supporting scalable and secure overlay network. The selective group of switches in EVPN multihoming mode shares common IPv4/IPv6 subnet with IP proxy function enable to support intra-subnet host communication based on fabric host-routing table while providing load-sharing and gateway redundancy with anycast gateway as routed overlay.

Spine route policy

The IPv4 or IPv6 subnet stretch across multiple IP gateway systems demand a non-condensed host-level prefix announcement in fabric core to support end-to-end seamless data communication. Switches that stretch the IPv4 or IPv6 subnets require the advertisement of the host-route to the fabric.

The iBGP or eBGP prefix advertisement to spine switches in fabric core limits advertising the IPv4 or IPv6 RT-2-MAC-IP host prefixes along with RT-5 network prefixes. The switches in EVPN multihoming mode retain unfiltered routing with iBGP peering session with the remote ES switch for non-blocking the Layer 2 multipath EtherChannel.



Note EVPN multihoming-enabled networks support additional overlay network types to stretch IP subnets or Layer 2 flood-domain beyond a single distribution block using Distributed Anycast Gateway (DAG) bridged technology. For better scale, performance, and resiliency, we recommend the routed overlay network.

Configure DAG routed overlay networks

This section provides step-by-step configuration guidelines to implement DAG routed overlay networks.

Configure IP VRFs

This task shows how to configure a logically segmented virtual network with an IP VRF to exchange IPv4 or IPv6 network prefixes across all VTEPs in the fabric core network.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vrf definition** *vrf-name* **Example:** `Device(config)# vrf definition green`	Configures a new virtual routing and forwarding (VRF) instance and enters VRF configuration mode.
Step 4	**rd** *vpn-route-distinguisher* **Example:** `Device(config-vrf)# rd 10.200.255.101:101`	Configures a route distinguisher with a distinct value across all fabric devices.
Step 5	**address-family ipv4** [**multicast** | **unicast**] **Example:** `Device(config-vrf-af)# address-family ipv4 unicast`	Enters IPv4 unicast address-family configuration mode.
Step 6	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **Example:** `Device(config- vrf-af)# route-target 65101:101`	Configures an overlay IPv4 import and export routing policy that matches the extended community value. • The default is **both** and will auto-generate both import and export command line with the configured community value.
Step 7	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **stitching** **Example:** `Device(config- vrf-af)# route-target 65101:101 stitching`	Configures an IPv4 routing import and export overlay policy that matches the extended community value and enables extended NLRI types for BGP EVPN VXLAN-enabled networks. • The default is **both**
Step 8	**address-family ipv6** [**multicast** | **unicast**] **Example:** `Device(config- vrf-af)# address-family ipv6 unicast`	Enters IPv6 unicast address-family configuration mode.

	Command or Action	Purpose
Step 9	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **Example:** `Device(config- vrf-af)# route-target 65101:101`	Configures an IPv6 routing import and export overlay policy that matches the extended community value. • The default is **both** and will auto-generate both import and export command line with the configured community value.
Step 10	**route-target** [**both** | **export** | **import**] *route-target-ext-community* **stitching** **Example:** `Device(config- vrf-af)# route-target 65101:101 stitching`	Configures an IPv6 routing import and export overlay policy that matches the extended community value and enables extended NLRI types for BGP EVPN VXLAN-enabled networks. • The default is **both**.
Step 11	**end** **Example:** `Device(config- vrf-af)# end`	Exits IPv6 unicast address-family configuration mode and returns to privileged EXEC mode.

Configure core VLAN and SVI interfaces

This task shows how to configure core VLAN and SVI interfaces for each of the IP VRFs used for VXLAN encapsulation and decapsulation functions between the downstream non-VXLAN network and remote VXLAN-enabled VTEPs. The task also shows how to configure the NVE interface and bind the L3VNI ID to a logical interface.

Note Skip this task if you have already configured core VLAN and SVI interfaces as part of the routed overlay network configuration.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**vlan** *id* **Example:**	Configures a dedicated VLAN ID for IP VRF core routing. This VLAN must not be used for

	Command or Action	Purpose
	`Device(config)# vlan 101`	any other purpose, and it is recommended to be pruned from Layer 2 trunk ports.
Step 4	**vlan configuration** *id* **Example:** `Device(config)# vlan configuration 101`	Configures a core VLAN to assign the Layer 3 VNI ID and enters VLAN configuration mode.
Step 5	**member vni** *l3vni-id* **Example:** `Device(config-vlan)# member vni 11101`	Configures a unique Layer 3 VNI ID for a specific IP VRF. For example, VRF green has the dedicated VLAN 101 and Layer 3 VNI, and these IDs must not be shared with other IP VRFs on the same system. All other VTEPs are recommended to have a Layer 3 VNI ID for the same IP VRF.
Step 6	**exit** **Example:** `Device(config-vlan)# exit`	Exits VLAN configuration mode and returns to global configuration mode
Step 7	**interface vlan** *id* **Example:** `Device(config)# interface vlan 101`	Configures a core VLAN SVI interface and assigns it to an IP VRF. The core VLAN SVI MAC address supports routed MAC (RMAC) for Layer 3 overlays.
Step 8	**vrf forwarding** *vrf-name* **Example:** `Device(config-if)# vrf forwarding green`	Maps the IP VRF to a core VLAN SVI interface.
Step 9	**ip unnumbered** *id* **Example:** `Device(config-if)# ip unnumbered Loopback0`	Configures an unnumbered IPv4 address from an underlay logical interface, for example, Loopback0.
Step 10	**no autostate** **Example:** `Device(config-if)# no autostate`	Disables autostate on the interface. • In EVPN deployments, a VLAN used for a core-facing SVI, should not be allowed in any trunks. For a core-facing SVI to function correctly, the **no autostate** command must be configured for the SVI.
Step 11	**exit** **Example:** `Device(config-if)# exit`	Exits interface configuration mode and returns to global configuration mode.

	Command or Action	Purpose
Step 12	**interface nve** *id* **Example:** `Device(config)# interface nve`	Configures a virtual NVE interface to bind one or more L2VNIs and enables system-wide VXLAN forwarding function from the L3VNIs. • Enters interface configuration mode.
Step 13	**source-interface** *id* **Example:** `Device(config-if)# source-interface Loopback0`	Configures a Loopback interface as the source to enable data communication over a VXLAN tunnel.
Step 14	**host-reachability protocol bgp** **Example:** `Device(config-if)# host-reachability protocol bgp`	Configures BGP as the host-reachability control plane protocol on this interface.
Step 15	**member vni** *l3vni-id* **vrf** *vrf-name* **Example:** `Device(config-if)# member vni 11101 vrf green`	Binds the L3VNI to the IP VRF under the NVE interface.
Step 16	**end** **Example:** `Device(config-if)# end`	Exits interface configuration mode and returns to privileged EXEC mode.

Configure DAG routed MAC VRFs

This task shows how to configure a MAC VRF and route-target policy to selectively exchange prefixes between a pair of EVPN multihoming-enabled switches in a single cluster.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**l2vpn evpn instance** *id* **vlan-based** **Example:**	Configures a L2VPN EVPN virtual instance with a unique ID. Typically, the ID value is the

	Command or Action	Purpose
	`Device(config)# l2vpn evpn instance 111 vlan-based`	same as the VLAN ID, for example, data VLAN ID 111. • Enters EVPN instance configuration mode.
Step 4	**encapsulation vxlan** **Example:** `Device(config-evpn-evi)# encapsulation vxlan`	Enables VXLAN encapsulation to bridge inter-ES data between two systems.
Step 5	**route-target** {**export** | **import** | **both**} *route-target-ext-community* **Example:** `Device(config-evpn-evi)# route-target both 1.1.1.1:111`	Configures a MAC or IP routing import and export policy that matches the extended community value. • The recommended format for EVPN multihoming cluster VLAN ID is Cluster-1 ID with 1.1.1.1 and DAG-routed VLAN ID 111 can result in the MAC VRF route-target 1.1.1.1:111. The default is **both** and it auto-generates the import and export command-line with the configured community value.
Step 6	**no auto-route-target** **Example:** `Device(config-evpn-evi)# no auto-route-target`	Disables auto route-target with the manual route-target configured in Step-5 for the DAG-routed overlay MAC VRF. • The default is auto-generated route-target in ASN:L2VPN instance ID format.
Step 7	**end** **Example:** `Device end`	Exits EVPN instance configuration mode and returns to privileged EXEC mode.

Configure IPv4 ARP and IPv6 ND/DAD proxy

Configures IPv4 Address Resolution Protocol (ARP) and IPv6 Neighbor Discovery (ND) or Duplicate Address Detection (DAD) proxy on switches to provide local anycast interface MAC response after receiving requests or queries from local endpoints connected to the downstream Layer 2 network.

This configuration is required only on access-facing VLAN or SVI interfaces that are implemented in distributed anycast gateway routed overlay mode

Procedure

	Command or Action	Purpose
Step 1	**enable**	Enables privileged EXEC mode.

	Command or Action	Purpose
	Example: `Device> enable`	• Enter your password if prompted
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**interface vlan** *id* **Example:** `Device(config)# interface vlan 111`	Configures a Layer 3 SVI interface ID and enters interface configuration mode.
Step 4	**ip address** *address mask* **Example:** `Device(config-if)# ip address 10.111.1.254 255.255.255.0`	Assigns IPv4 gateway address and mask for selected VLANs. • Anycast gateway address is auto enabled with a common address between ES systems.
Step 5	**ip local-proxy-arp** **Example:** `Device(config-if)# ip local-proxy-arp`	Configures a gateway to send proxy responses with the local interface MAC address when receiving IPv4 ARP requests from local hosts.
Step 6	**ipv6 address** *address mask* **Example:** `Device(config-if)# ipv6 address 2001:10:111:1::f/64`	Assigns IPv6 gateway address and prefix for selected VLANs. • Anycast gateway address is auto-enabled with a common address between ES systems.
Step 7	**exit** **Example:** `Device(config-if)# exit`	Exits interfaced configuration mode and returns to privileged EXEC mode.
Step 8	**vlan configuration** *id* **Example:** `Device(config)# vlan configuration 111`	Enters VLAN configuration mode
Step 9	**ipv6 nd routing proxy** **Example:** `Device(config-vlan)# ipv6 nd routing proxy`	Configures a gateway switch to send proxy responses with the local interface MAC address when receiving IPv6 ND requests from local hosts.
Step 10	**ipv6 nd dad-proxy** **Example:** `Device(config-vlan)# ipv6 nd dad-proxy`	Configures the gateway switch to send DAD proxy response to local hosts with the local interface MAC address when receiving IPv6 ND DAD messages for the IPv6 address owned by another host.

	Command or Action	Purpose
Step 11	**end** **Example:** `Device(config-vlan)# end`	Exits VLAN configuration mode and returns to privileged EXEC mode.

Configure spine policies for DAG routed overlay networks

This task provides step-by-step instructions to configure a pair of ES switches with an outbound spine EVPN route policy. The route policy enables controlled EVPN route advertisement by the pair of ES switches, independent of iBGP or eBGP spine-peering session type.

The spine route-map policy to permit RT-5 prefix for routed overlay can be expanded with a new sequence to limit the switches advertising the RT-2 with combined MAC and IP (RT-2-MAC-IP) prefix for network-wide host route with IPv4 or IPv6 address propagation.

The iBGP peering with EVPN multihoming system unfiltered EVPN route exchange to build and maintain non-blocking Layer 2 multipath network.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**route-map** *name* [**permit** | **deny**] *sequence-number* **Example:** `Device(config)# route-map SPINE-ROUTE-POLICY-OUT permit 20`	Creates a route-map rule with permit or deny match criteria and enters route-map configuration mode. • The optional *sequence-number* argument indicates the order of processing of the route-map rule.
Step 4	**description** *description* **Example:** `Device(config-route-map)# description DAG ROUTED OVERLAY – VL111`	(Optional) Adds a description for the route-map sequence.
Step 5	**match evpn route-type** [**1-8** | **2-mac-ip** | **2-mac-only**] **Example:** `Device(config-route-map)# match evpn route-type 2-mac-ip`	Configures an EVPN route-policy sequence to permit advertising the local IPv4 or IPv6 hosts prefixes (RT-2-MAC-IP) mapped to any IP VRF.

	Command or Action	Purpose
Step 6	**exit** **Example:** `Device(config-route-map)# exit`	Exits route-map configuration mode and returns to global configuration mode.
Step 7	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 65101`	Configures BGP autonomous system numbers using 2-bytes or 4-bytes in asplain and asdot formats and enters router configuration mode.
Step 8	**address-family l2vpn evpn** **Example:** `Device(config-router)# address-family l2vpn evpn`	Enters BGP L2VPN address-family configuration mode.
Step 9	**neighbor** *spine-loopback-address* **route-map** *name* **out** **Example:** `Device(config-router-af)# neighbor 10.200.255.3 route-map SPINE-ROUTE-POLICY-OUT out`	Applies outbound route-map policy to each spine peer to limit the advertising of the RT-5 prefix.
Step 10	**end** **Example:** `Device(config-router-af)# end`	Exits router configuration mode and returns to privileged EXEC mode.

Configure IP VRF address family routing under BGP

This task shows how to advertise network prefixes in IPv4 and IPv6 address-family configuration.

Procedure

	Command or Action	Purpose
Step 1	**enable** **Example:** `Device> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Device# configure terminal`	Enters global configuration mode.
Step 3	**router bgp** *autonomous-system-number* **Example:** `Device(config)# router bgp 65101`	Configures BGP autonomous system numbers using 2-bytes or 4-bytes in asplain and asdot formats and enters router configuration mode.

	Command or Action	Purpose
Step 4	**address-family ipv4 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv4 vrf green`	Enables IPv4 address-family for a VRF instance to enable IPv4 network prefix routing in overlay.
Step 5	**advertise l2vpn evpn** **Example:** `Device(config-route-af)# advertise l2vpn evpn`	Enables advertising the IPv4 network prefix routing between IPv4 VRF instances and EVPN fabric core networks.
Step 6	**redistribute connected** [**metric** |**route-map**] *name* **Example:** `Device(config-route-af)# redistribute connected`	Configures IPv4 address-family to advertise all local connected IPv4 network prefixes in EVPN fabric network. • The network prefix can be advertised with network command-line as an alternative.
Step 7	**exit** **Example:** `Device(config-route-af)# exit`	Exits IPv4 address-family configuration mode and returns to router configuration mode
Step 8	**address-family ipv6 vrf** *vrf-name* **Example:** `Device(config-router)# address-family ipv6 vrf green`	Enables IPv6 address-family for a VRF instance to enable IPv6 network prefix routing in overlay.
Step 9	**advertise l2vpn evpn** **Example:** `Device(config-route-af)# advertise l2vpn evpn`	Enables IPv6 network prefix routing between IPv6 VRF instances and EVPN fabric core networks.
Step 10	**redistribute connected** [**metric** |**route-map**] *name* **Example:** `Device(config-route-af)# redistribute connected`	Configures an IPv6 address-family to advertise all local connected IPv6 network prefixes in EVPN fabric network. • The network prefix can be advertised with network command-line as an alternative.
Step 11	**end** **Example:** `Device(config-route-af)# end`	Exits IPv6 address-family configuration mode and returns to privileged EXEC mode.

Verify the DAG routed overlay configuration

This section provides examples to verify the DAG routed overlay network configuration and states of VTEPs.

The command output may be truncated to focus on critical information for a Day-2 operation and troubleshooting.

IP VRF networks: Verifies the locally configured IP VRF and the associated physical or logical interface bindings to a virtual network on VTEPs including leaf and border devices. The interface binding includes the network edge VLANs connecting to endpoints, and network devices, along with a single core VLAN (VLAN 101) of an IP VRF.

```
ES-1# show vrf green

  Name      Default RD            Protocols   Interfaces
  green     10.200.255.101:101    ipv4   Vl11
                                             Vl111
                                             Vl101

BORDER-1# show vrf green

  Name      Default RD            Protocols    Interfaces
  green     10.200.255.1:101      ipv4    Vl101
                                              Vl1101
```

BGP L2VPN EVPN neighbors: Verifies that the two-tier hierarchical iBGP sessions between switches in EVPN multihoming mode and iBGP peering to a pair of spine systems are in operational state.

The following command output shows iBGP peering with a pair of spines switches, 10.100.255.3 and 10.100.255.4 and direct iBGP peering between ES-1, local 10.100.255.101 and ES-2, 10.100.255.102 switches in EVPN multihoming mode is operational.

```
ES-1# show bgp l2vpn evpn all summary

BGP router identifier 10.100.255.101, local AS number 65101
<snip>

Neighbor        V    AS     MsgRcvd  MsgSent   TblVer  InQ   OutQ    Up/Down    State/PfxRcd

10.100.255.3     4    65101    18        20      104     0      0     00:04:35         2
10.100.255.4     4    65101    23        25      106     0      0     00:05:19         2
10.100.255.102   4    65101    51        65      104     0      0     00:04:26         28
```

Spine policy: Verifies the VTEP in leaf or border roles implemented in EVPN multihoming mode is configured with a route-map and an outbound policy is applied to each spine, 10.200.255.3.

The spine policy configuration on border switches is optional and not required if the switches are not implemented in EVPN multihoming mode to connect to external network devices, such as firewalls.

```
ES-1# show bgp l2vpn evpn neighbor 10.200.255.3 policy

Neighbor: 10.200.255.3, Address-Family: L2VPN E-VPN
 Locally configured policies:
  send-community extended
 Inherited polices:
  route-map EVPN-SPINE-ROUTE-POLICY-OUT out
```

L2VPN advertised routes: Verifies the VTEP in leaf or border roles implemented in EVPN multihoming mode is advertising the EVPN prefixes to the spine, 10.200.255.3, based on the applied route-map policy.

The following command output from leaf and border switches confirms the advertising of RT-2-MAC-IP EVPN prefixes, and that no additional EVPN route-types are advertised to the spine switches.

```
ES-1# show bgp l2vpn evpn neighbor 10.200.255.3 advertised-route

BGP table version is 7179, local router ID is 10.200.255.101
<snip>
     Network          Next Hop            Metric  LocPrf  Weight  Path
 Route Distinguisher: 10.200.255.101:111
 *>   [2][10.200.255.101:111][0][48][00CCFC4ED1C3][32][10.111.1.1]/24
                      0.0.0.0                      32768     ?
Route Distinguisher: 10.200.255.101:111
 *>   [2][10.200.255.101:111][0][48][00CCFC4ED1C4][32][10.111.1.2]/24
                      0.0.0.0                           32768     ?


Border-1# show bgp l2vpn evpn neighbors 10.200.255.3 advertised-routes

BGP table version is 2895, local router ID is 10.200.255.1
<snip>
     Network          Next Hop            Metric    LocPrf     Weight   Path
Route Distinguisher: 10.100.255.1:101 (default for vrf green)
 *>   [5][10.100.255.1:101][0][0][0.0.0.0]/17
                      21.1.1.1               0          0      65001      ?
 *>   [5][10.100.255.1:101][0][16][111.1.0.0]/17
                      21.1.1.1               0          0      65001      ?
```

MAC-VRF manual route-target: Verifies that the stretched IP-subnet VLAN is limited to import and export of the MAC or IP route matching the local cluster EVPN multihoming peer switch. With auto-route-target disabled, the import and export route-target only lists the manual configuration settings.

The following command output confirms that the bidirectional MAC route policy with route-target 1.1.1.1:111 and the default auto-generated ASN:Instance ID 65101:111 is not listed along with the manual configuration.

```
ES-1# show l2vpn evpn evi 111 detail

EVPN instance:           111 (VLAN Based)
  RD:                    10.200.255.101:111 (auto)
  Import-RTs:            1.1.1.1:111
  Export-RTs:            1.1.1.1:111
  Per-EVI Label:         none
  State:                 Established
  Replication Type:      Ingress (global)
 <snip>
```

IP local proxy: Verifies the stretched IP-subnet SVI interface configured to support the IPv4 ARP local proxy function.

Use the **show ip interface vlan** command to verify the Anycast Gateway SVI interface. Network administrators can verify the ARP entries on local hosts to confirm the common IP gateway and the remote host IP-to-MAC binding entries.

```
ES-1# show ip interface vlan111

Vlan111 is up, line protocol is up
  Internet address is 10.111.1.254/24
  Broadcast address is 255.255.255.255
  Address determined by non-volatile memory
```

```
  MTU is 9100 bytes
 <snip>
  Proxy ARP is enabled
  Local Proxy ARP is enabled
<snip>
```

VRF IP routing table: Verifies that the leaf implemented in EVPN multihoming mode is advertising the EVPN prefixes to the spine switch 10.200.255.3 based on the applied route-map policy.

The following command output from the border switch confirms that the pair of EVPN multihoming leaf systems configured in DAG routed overlay, each with two hosts are advertising RT2-MAC-IP. Border switches have a dual-path ECMP for each host to support load-sharing and redundancy for overlay network prefixes. And leaf switches receive RT-5 network prefix from the IP routing table of IP VRFs in the border systems that confirm the advertising of RT-5 EVPN prefixes.

```
ES-1# show ip route vrf green bgp

Routing Table: green
<snip>

Gateway of last resort is 10.200.255.2 to network 0.0.0.0
B*    0.0.0.0/0  [200/0] via 10.200.255.1, 00:00:32, Vlan101
[200/0] via 10.200.255.2, 00:00:32, Vlan101


Border-1# show ip route vrf green bgp
Routing Table: green
<snip>
Gateway of last resort is 21.1.1.1 to network 0.0.0.0

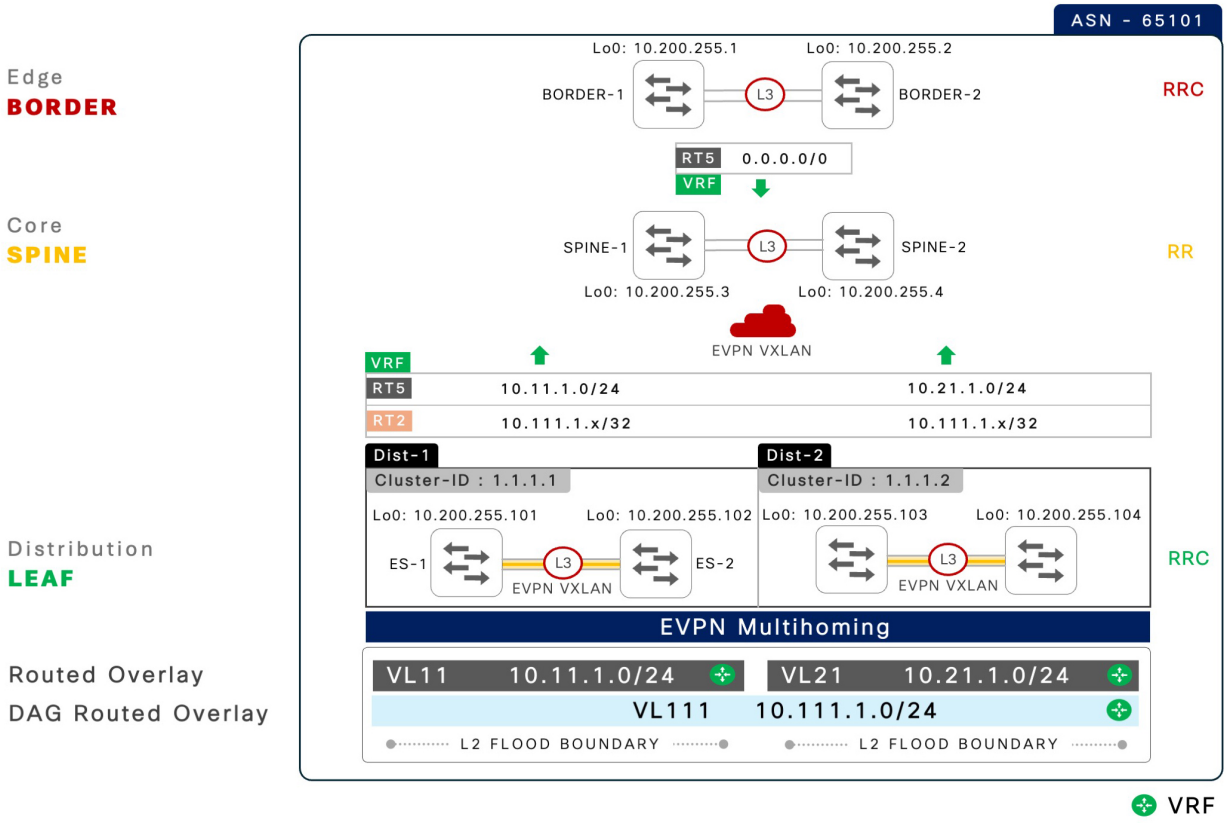
B*    0.0.0.0/0 [20/0] via 21.1.1.1, 6w0d
      10.0.0.0/24 is subnetted, 95 subnets
B        10.111.1.1/32 [200/0] via 10.200.255.102, 00:01:33, Vlan1101
                       [200/0] via 10.200.255.101, 00:01:33, Vlan1101
B        10.111.1.2/32 [200/0] via 10.200.255.102, 00:01:33, Vlan1101
                           [200/0] via 10.200.255.101, 00:01:33, Vlan1101
B        10.111.1.3/32  [200/0] via 10.200.255.104, 00:01:33, Vlan1101
                           [200/0] via 10.200.255.103, 00:01:33, Vlan1101
B        10.111.1.4/32  [200/0] via 10.200.255.104, 00:01:33, Vlan1101
                           [200/0] via 10.200.255.103, 00:01:33, Vlan1101
```

Reference configuration for DAG routed overlay networks

This section provides a reference network design and configuration to implement end-to-end Layer 2 flood-free DAG routed overlay fabric network.

This following figure illustrates a reference network design to support scalable EVPN multihoming and fabric core with hierarchical BGP peering, control-plane peering, and selective stretching of an IP subnet between a pair of switches deployed in the distribution block with overlay fabric network.

Figure 78: EVPN multihoming: DAG routed overlay fabric network



This reference configuration shows a three-step process, from the initial network setup to the successful fabric deployment. Network administrators can follow these steps sequentially for initial deployment and to increment the overlay network configuration as the network demands increase.

Layer 2 EVPN multihoming

The base configuration to build a fabric begins with Layer 2 campus networks using EVPN multihoming technology. This section provides the step-by-step configuration on a pair of switches to successfully build a non-blocking all-active Layer 2 network with EVPN multihoming.

Step	ES-1 and ES-2	ES-3 and ES-4
1: Inter-ES Layer 3 EtherChannel	ES-1 ! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-2 no switchport ip address 10.0.0.0 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 101 ip ospf cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking ! ES-2 ! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-2 no switchport ip address 10.0.0.1 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 101 ip ospf cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !	ES-3 ! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-1 no switchport ip address 10.0.0.2 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 102 ip ospf cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking ! ES-4 ! interface Port-Channel 128 description CONNECTED TO EVPN MH SW-1 no switchport ip address 10.0.0.3 255.255.255.254 ip ospf network point-to-point ip ospf multi-area 0 ip ospf 100 area 102 ip ospf cost 10 carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !

Step	ES-1 and ES-2	ES-3 and ES-4

2: IGP routing and core interface	ES-1
	!
	router ospf 100
	 router-id 10.200.255.101
	 max-metric router-lsa include-stub summary-lsa
	  external-lsa on-startup wait-for-bgp
	
	 nsf cisco
	 fast-reroute per-prefix enable prefix-priority
	  low
	 area 101 stub no-summary
	 passive-interface default
	 no passive-interface Port-Channel 128
	
	 no passive-interface HundredGig1/0/49
	
	 no passive-interface HundredGig1/0/50
	
	!
	interface Loopback 0
	 ip address 10.100.255.101 255.255.255.255
	 ip ospf 100 area 0
	!
	ES-2
	!
	router ospf 100
	 router-id 10.200.255.102
	 max-metric router-lsa include-stub summary-lsa
	  external-lsa on-startup wait-for-bgp
	
	 nsf cisco
	 fast-reroute per-prefix enable prefix-priority
	  low
	 area 101 stub no-summary
	 passive-interface default
	 no passive-interface Port-Channel 128
	
	 no passive-interface HundredGig1/0/49
	
	 no passive-interface HundredGig1/0/50
	
	!
	interface Loopback 0
	 ip address 10.100.255.102 255.255.255.255
	 ip ospf 100 area 0
	!
		ES-3
		!
		router ospf 100
		 router-id 10.200.255.103
		 max-metric router-lsa include-stub summary-lsa
		  external-lsa on-startup wait-for-bgp
		 nsf cisco
		 fast-reroute per-prefix enable prefix-priority
		  low
		 area 102 stub no-summary
		 passive-interface default
		 no passive-interface Port-Channel 128
		 no passive-interface HundredGig1/0/49
		 no passive-interface HundredGig1/0/50
		!
		interface Loopback 0
		 ip address 10.100.255.103 255.255.255.255
		 ip ospf 100 area 0
		!
		ES-4
		!
		router ospf 100
		 router-id 10.200.255.104
		 max-metric router-lsa include-stub summary-lsa
		  external-lsa on-startup wait-for-bgp
		 nsf cisco
		 fast-reroute per-prefix enable prefix-priority
		  low
		 area 102 stub no-summary
		 passive-interface default
		 no passive-interface Port-Channel 128
		 no passive-interface HundredGig1/0/49
		 no passive-interface HundredGig1/0/50
		!
		interface Loopback 0
		 ip address 10.100.255.104 255.255.255.255
		 ip ospf 100 area 0
		!

Step	ES-1 and ES-2	ES-3 and ES-4
3: iBGP routing		

Step	ES-1 and ES-2	ES-3 and ES-4

	ES-1
	!
	router bgp 65101
	 template peer-policy ES-PEER-POLICY
	  send-community both
	 !
	 template peer-session
	ES-PEER-SESSION-POLICY
	  remote-as 65101
	  description EVPN-MH-DIST-2-PEER
	  update-source Loopback0
	  fall-over host-route
	 !
	 bgp router-id interface Loopback0
	 bgp log-neighbor-changes
	 bgp graceful-restart
	 no bgp default ipv4-unicast
	 neighbor 10.100.255.102 inherit
	peer-session
	  ES-PEER-SESSION-POLICY
	 !
	 address-family l2vpn evpn
	  bgp nexthop trigger critical-delay 0
	
	  neighbor 10.100.255.102 activate
	  neighbor 10.100.255.102 send-community
	 both
	  neighbor 10.100.255.102 inherit
	peer-policy
	  ES-PEER-POLICY
	 !
	ES-2
	!
	router bgp 65101
	 template peer-policy ES-PEER-POLICY
	  send-community both
	 !
	 template peer-session
	ES-PEER-SESSION-POLICY
	  remote-as 65101
	  description EVPN-MH-DIST-2-PEER
	  update-source Loopback0
	  fall-over host-route
	 !
	 bgp router-id interface Loopback0
	 bgp log-neighbor-changes
	 bgp graceful-restart
	 no bgp default ipv4-unicast
	 neighbor 10.100.255.101 inherit
	peer-session
	  ES-PEER-SESSION-POLICY
	 !
	 address-family l2vpn evpn
	  bgp nexthop trigger critical-delay 0
	
	  neighbor 10.100.255.101 activate
	  neighbor 10.100.255.101 send-community
	 both
	  neighbor 10.100.255.101 inherit
	peer-policy
	  ES-PEER-POLICY
	 !
		ES-3
		!
		router bgp 65101
		 template peer-policy ES-PEER-POLICY
		
		  send-community both
		 !
		 template peer-session
		ES-PEER-SESSION-POLICY
		  remote-as 65101
		  description EVPN-MH-DIST-1-PEER
		  update-source Loopback0
		  fall-over host-route
		 !
		 bgp router-id interface Loopback0
		 bgp log-neighbor-changes
		 bgp graceful-restart
		 no bgp default ipv4-unicast
		 neighbor 10.100.255.104 inherit
		peer-session
		  ES-PEER-SESSION-POLICY
		 !
		 address-family l2vpn evpn
		  bgp nexthop trigger critical-delay
		 0
		  neighbor 10.100.255.104 activate
		  neighbor 10.100.255.104
		send-community both
		  neighbor 10.100.255.104 inherit
		peer-policy
		  ES-PEER-POLICY
		 !
		ES-4
		!
		router bgp 65101
		 template peer-policy ES-PEER-POLICY
		
		  send-community both
		 !
		 template peer-session
		ES-PEER-SESSION-POLICY
		  remote-as 65101
		  description EVPN-MH-DIST-1-PEER
		  update-source Loopback0
		  fall-over host-route
		 !
		 bgp router-id interface Loopback0
		 bgp log-neighbor-changes
		 bgp graceful-restart
		 no bgp default ipv4-unicast
		 neighbor 10.100.255.103 inherit
		peer-session
		  ES-PEER-SESSION-POLICY
		 !
		 address-family l2vpn evpn
		  bgp nexthop trigger critical-delay
		 0
		  neighbor 10.100.255.103 activate
		  neighbor 10.100.255.103
		send-community both
		  neighbor 10.100.255.103 inherit
		peer-policy

Step	ES-1 and ES-2	ES-3 and ES-4
		ES-PEER-POLICY !
4: Global L2VPN	! l2vpn evpn advertise mac disable anycast-gateway mac auto multicast advertise sync-only multihoming aliasing disable multihoming peering adjacent replication-type ingress router-id Loopback 0 !	! l2vpn evpn advertise mac disable anycast-gateway mac auto multicast advertise sync-only multihoming aliasing disable multihoming peering adjacent replication-type ingress router-id Loopback 0 !
5: Routed VLAN and MAC VRF	! vlan 11 name ROUTED_DATA_VLAN ! l2vpn evpn instance 11 vlan-based encapsulation vxlan ! vlan configuration 11 member evpn-instance 11 vni 11011 ! interface nve 1 source-interface Loopback 0 host-reachability protocol bgp member vni 11011 ingress-replication !	! vlan 11 name ROUTED_DATA_VLAN ! l2vpn evpn instance 11 vlan-based encapsulation vxlan ! vlan configuration 11 member evpn-instance 11 vni 11011 ! interface nve 1 source-interface Loopback 0 host-reachability protocol bgp member vni 11011 ingress-replication !
6: DAG-routed VLAN and MAC VRF	! vlan 111 name DAG_ROUTED_DATA_VLAN ! l2vpn evpn instance 111 vlan-based encapsulation vxlan route-target 1.1.1.1:111 no auto-route-target ! vlan configuration 111 member evpn-instance 111 vni 11111 ! interface nve 1 member vni 11111 ingress-replication !	! vlan 111 name DAG_ROUTED_DATA_VLAN ! l2vpn evpn instance 111 vlan-based encapsulation vxlan route-target 1.1.1.1:111 no auto-route-target ! vlan configuration 111 member evpn-instance 111 vni 11111 ! interface nve 1 member vni 11111 ingress-replication !
7: ES EtherChannel	! interface Port-Channel 1 description CONNECTED TO L2 ACCESS switchport trunk allowed vlan 11,111 evpn ethernet-segment auto lacp df-election wait-time 1 !	! interface Port-Channel 1 description CONNECTED TO L2 ACCESS switchport trunk allowed vlan 11,111 evpn ethernet-segment auto lacp df-election wait-time 1 !

Underlay: fabric core and BGP peering

Enterprise campus core networks with solid underlay network foundation are the key for highly scalable, resilient BGP EVPN VXLAN fabric networks. This section is the second step to build a reliable underlay core network for fabric and hierarchical BGP peering on targeted network devices with specific roles.



Note The table is subdivided into two fabric roles with each step either sharing common configuration or a unique per-device with a common role.

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
1: Global best practices	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !	! system mtu 9100 ! port-channel load-balance vlan-src-dst-mixed-ip-port ip cef load-sharing algorithm include-ports source destination protocol ! ip tcp mss 8000 ip tcp window-size 262144 ip tcp path-mtu-discovery !
2: Underlay interface configuration and best practices	! interface range HundredGig1/0/49-50 description CONNECTED TO SPINE ip ospf 100 area 0 ip ospf network point-to-point carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out evpn multihoming core-tracking !	! interface range HundredGig1/0/1-4 description CONNECTED TO CAMPUS CORE NETWORK ip ospf 100 area 0 ip ospf network point-to-point carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out !	! interface range HundredGig1/0/49-50 description CONNECTED TO SPINE ip ospf 100 area 0 ip ospf network point-to-point carrier-delay msec 0 hold-queue 4094 in hold-queue 4094 out !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
3: OSPF routing configuration and best practices	ES-1 and ES-2 ! router ospf 100 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low area 101 stub no-summary passive-interface default no passive-interface Port-Channel 128 no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 ! ES-3 and ES-4 ! router ospf 100 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low area 102 stub no-summary passive-interface default no passive-interface Port-Channel 128 no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 !	SPINE-1 ! router ospf 100 router-id 10.200.255.3 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/1 no passive-interface HundredGig1/0/2 no passive-interface HundredGig1/0/3 no passive-interface HundredGig1/0/4 ! SPINE-2 ! router ospf 100 router-id 10.200.255.4 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/1 no passive-interface HundredGig1/0/2 no passive-interface HundredGig1/0/3 no passive-interface HundredGig1/0/4 !	BORDER-1 ! router ospf 100 router-id 10.200.255.1 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 ! BORDER-2 ! router ospf 100 router-id 10.200.255.2 max-metric router-lsa include-stub summary-lsa external-lsa on-startup wait-for-bgp nsf cisco fast-reroute per-prefix enable prefix-priority low passive-interface default no passive-interface HundredGig1/0/49 no passive-interface HundredGig1/0/50 !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
4:BGP routing configuration and best practices	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !	! router bgp 65101 ! bgp router-id interface Loopback0 bgp log-neighbor-changes bgp graceful-restart no bgp default ipv4-unicast !
5: Peer-session and peer-policy templates and parameters for leaf switches	! template peer-session EVPN-SPINE-PEER-SESSION-POLICY remote-as 65101 description EVPN-SPINE-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-SPINE-PEER-POLICY send-community both !	! template peer-session EVPN-LEAF-PEER-SESSION-POLICY remote-as 65101 description EVPN-LEAF-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-LEAF-PEER-POLICY route-reflector-client send-community both !	
6: Peer-session and policy templates and parameters for border switches		! template peer-session EVPN-BORDER-PEER-SESSION-POLICY remote-as 65101 description EVPN-BORDER-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-BORDER-PEER-POLICY route-reflector-client send-community both !	! template peer-session EVPN-SPINE-PEER-SESSION-POLICY remote-as 65101 description EVPN-SPINE-PEER log-neighbor-changes update-source Loopback0 fall-over host-route ! template peer-policy EVPN-SPINE-PEER-POLICY send-community both !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
7: Disable intra-cluster EVPN multihome leaf reflection		! no bgp client-to-client reflection intra-cluster cluster-id any	
8: Border-spine iBGP peering		! neighbor 10.200.255.1 inherit peer-session EVPN-BORDER-PEER-SESSION-POLICY ! neighbor 10.200.255.2 inherit peer-session EVPN-BORDER-PEER-SESSION-POLICY !	! neighbor 10.200.255.3 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY ! neighbor 10.200.255.4 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
9: Leaf iBGP peering	! neighbor 10.200.255.3 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY ! neighbor 10.200.255.4 inherit peer-session EVPN-SPINE-PEER-SESSION-POLICY !	! neighbor 10.200.255.101 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.101 cluster-id 1.1.1.1 ! neighbor 10.200.255.102 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.102 cluster-id 1.1.1.1 ! neighbor 10.200.255.103 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.101 cluster-id 1.1.1.2 ! neighbor 10.200.255.104 inherit peer-session EVPN-LEAF-PEER-SESSION-POLICY neighbor 10.200.255.102 cluster-id 1.1.1.2 !	

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
10: Activate leaf and border iBGP peering under L2VPN EVPN address family	! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.3 activate neighbor 10.200.255.3 send-community both neighbor 10.200.255.3 inherit peer-policy EVPN-SPINE-PEER-POLICY neighbor 10.200.255.4 activate neighbor 10.200.255.4 send-community both neighbor 10.200.255.4 inherit peer-policy EVPN-SPINE-PEER-POLICY !		! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.3 activate neighbor 10.200.255.3 send-community both neighbor 10.200.255.3 inherit peer-policy EVPN-SPINE-PEER-POLICY neighbor 10.200.255.4 activate neighbor 10.200.255.4 send-community both neighbor 10.200.255.4 inherit peer-policy EVPN-SPINE-PEER-POLICY !

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
		! address-family l2vpn evpn bgp nexthop trigger critical-delay 0 neighbor 10.200.255.1 activate neighbor 10.200.255.1 send-community both neighbor 10.200.255.1 inherit peer-policy EVPN-BORDER-PEER-POLICY neighbor 10.200.255.2 activate neighbor 10.200.255.2 send-community both neighbor 10.200.255.2 inherit peer-policy EVPN-BORDER-PEER-POLICY ! neighbor 10.200.255.101 activate neighbor 10.200.255.101 send-community both neighbor 10.200.255.101 inherit peer-policy EVPN-LEAF-PEER-POLICY neighbor 10.200.255.102 activate neighbor 10.200.255.102 send-community both neighbor 10.200.255.102 inherit peer-policy EVPN-LEAF-PEER-POLICY ! neighbor 10.200.255.103 activate neighbor 10.200.255.103 send-community both	

Step	ES-1, ES-2, ES-3 and ES-4	Spine-1 and Spine-2	Border-1 and Border-2
		neighbor 10.200.255.103 inherit peer-policy EVPN-LEAF-PEER-POLICY neighbor 10.200.255.104 activate neighbor 10.200.255.104 send-community both neighbor 10.200.255.104 inherit peer-policy EVPN-LEAF-PEER-POLICY !	

Overlay: DAG routed networks

The overlay network configuration is the final step to enable a fabric in the enterprise campus. This section provides step-by-step configuration procedures on VTEPs involved in a DAG routed overlay network that exchanges IP prefixes between external and internal network domains.

Step	ES-1, ES-2, ES-3 and ES-4	Border-1 and Border-2
1: IP VRF configuration	ES-1 ! vrf definition green rd 10.200.255.101:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-2 ! vrf definition green rd 10.200.255.102:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-3 ! vrf definition green rd 10.200.255.103:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! ES-4 ! vrf definition green rd 10.200.255.104:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching !	BORDER-1 ! vrf definition green rd 10.200.255.1:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching ! BORDER-2 ! vrf definition green rd 10.200.255.2:101 address-family ipv4 unicast route-target 65101:101 route-target 65101:101 stitching !
2: IP VRF core VLAN configuration	! vlan 101 name VRF_GREEN_CORE_VLAN ! vlan configuration 101 member vni 10011 ! interface vlan 101 description CORE VLAN - VRF GREEN vrf forwarding green ip unnumbered Loopback0 no autostate !	! vlan 101 name VRF_GREEN_CORE_VLAN ! vlan configuration 101 member vni 10011 ! interface vlan 101 description CORE VLAN - VRF GREEN vrf forwarding green ip unnumbered Loopback0 no autostate !
3: IP VRF L3VNI to NVE interface binding	! interface nve 1 member vni 10011 vrf green !	! interface nve 1 member vni 10011 vrf green !

Step	ES-1, ES-2, ES-3 and ES-4	Border-1 and Border-2
4: Network edge to access or external domain.	ES-1 and ES-2 ! interface Vlan 11 description ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.11.1.254 255.255.255.0 ! interface Vlan 111 description DAG ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.111.1.254 255.255.255.0 ! ES-3 and ES-4 ! interface Vlan 21 description ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.21.1.254 255.255.255.0 ! interface Vlan 111 description DAG ROUTED DATA VLAN – VRF GREEN vrf forwarding green ip address 10.111.1.254 255.255.255.0 !	BORDER-1 ! interface Vlan 2001 description FIREWALL HANDOFF – VRF GREEN vrf forwarding green ip address 21.1.1.0 255.255.255.254 ! BORDER-2 ! interface Vlan 2002 description FIREWALL HANDOFF – VRF GREEN vrf forwarding green ip address 21.1.1.2 255.255.255.254 !
5: Route-map policy	! route-map SPINE-ROUTE-POLICY-OUT permit 10 description ROUTED OVERLAY NETWORK POLICY match evpn route-type 5 ! route-map SPINE-ROUTE-POLICY-OUT permit 20 description DAG ROUTED OVERLAY NETWORK POLICY match evpn route-type 2-mac-ip !	
6: Apply spine policy to BGP template	! router bgp 65101 ! template peer-policy EVPN-SPINE-PEER-POLICY route-map SPINE-ROUTE-POLICY-OUT out !	

Step	ES-1, ES-2, ES-3 and ES-4	Border-1 and Border-2
7: IP routing	! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected maximum-paths ibgp 2 !	BORDER-1 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.1 remote-as 65001 neighbor 21.1.1.1 activate maximum-paths ibgp 2 ! BORDER-2 ! router bgp 65101 ! address-family ipv4 vrf green advertise l2vpn evpn neighbor 21.1.1.3 remote-as 65001 neighbor 21.1.1.3 activate maximum-paths ibgp 2 !