



BGP EVPN VXLAN Configuration Guide, Cisco IOS XE Dublin 17.12.x (Catalyst 9300 Switches)

First Published: 2023-07-28

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 NON-INFRINGEMENT 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.

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

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)

© 2023 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

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
Example: Configuring BUM Traffic Rate Limiting	71

CHAPTER 3

Configuring EVPN VXLAN Layer 3 Overlay Network	75
Information About EVPN VXLAN Layer 3 Overlay Network	75
How to Configure EVPN VXLAN Layer 3 Overlay Network	76
Configuring an IP VRF on a VTEP	76
Configuring the Core-facing VLAN on a VTEP	78
Configuring Access-facing VLAN on a VTEP	79
Configuring Switch Virtual Interface for the Core-facing VLAN	79
Configuring the Switch Virtual Interface for the Access-facing VLANs	80
Configuring the Loopback Interface on a VTEP	81
Configuring the NVE Interface on a VTEP	82
Configuring BGP with IPv4 or IPv6 or Both Address Families on VTEP	83
Configuration Examples for EVPN VXLAN Layer 3 Overlay Network	86
Verifying EVPN VXLAN Layer 3 Overlay Network	93

CHAPTER 4	Configuring EVPN VXLAN Integrated Routing and Bridging	95
	Restrictions for EVPN VXLAN Integrated Routing and Bridging	95
	Information About EVPN VXLAN Integrated Routing and Bridging	95
	EVPN VXLAN Distributed Anycast Gateway	96
	Manual MAC Address Configuration	98
	MAC Aliasing	98
	EVPN VXLAN Centralized Default Gateway	98
	Default Gateway MAC Address Assignment	100
	Route Type 2 to Route Type 5 Host Route Reorigination for Overlay Networks with a Centralized Gateway (Asymmetric IRB)	100
	How to Configure EVPN VXLAN Integrated Routing and Bridging	101
	Configuring EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway	101
	Configuring Layer 2 VPN EVPN on a VTEP	101
	Configuring IP VRF on VTEP	101
	Configuring the Loopback Interface on a VTEP	101
	Configuring Core-facing and Access-facing VLANs on a VTEP	101
	Configuring Switch Virtual Interface for the Core-facing VLAN on a VTEP	102
	Configuring Switch Virtual Interface for the Access-facing VLANs on a VTEP	103
	Configuring the NVE Interface on a VTEP	104
	Configuring BGP with EVPN and VRF Address Families on a VTEP	105
	Configuring EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway	108
	Verifying EVPN VXLAN Integrated Routing and Bridging	109
	Configuration Examples for EVPN VXLAN Integrated Routing and Bridging	110
	Example: Enabling EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway	110
	Example: Enabling EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway	124
CHAPTER 5	Configuring Spine Switches in a BGP EVPN VXLAN Fabric	141
	Information About Spine Switches in a BGP EVPN VXLAN Fabric	141
	Deployment Scenarios for Spine Switches and Leaf Switches in a BGP EVPN VXLAN Fabric	141
	Configuration Examples for Spine Switches in a BGP EVPN VXLAN Network	142
	Configuration Example-Spine and Leaf Switches in the Same Autonomous System	142

Configuration Example-Spine and Leaf Switches are in Different Autonomous Systems 159
 Configuration Example-Each Leaf Switch in a Different Autonomous System 177

CHAPTER 6

Configuring DHCP Relay in a BGP EVPN VXLAN Fabric 197
 Information About DHCP Relay in a BGP EVPN VXLAN Fabric 197
 DHCP Relay on VTEPs 198
 How to Configure DHCP Relay in a BGP EVPN VXLAN Fabric 199
 Configuring DHCP Relay on a VTEP 199
 Configuring DHCP Relay on the Access SVI of a VTEP 200
 Configuring the Layer 3 or Routed Interface on the Border VTEP for DHCP Server Reachability 201
 Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric 203
 Example: DHCP Server is in the Layer 3 Default VRF and the DHCP Client is in the Tenant VRF 204
 Example: DHCP Server and DHCP Client are in the Same Tenant VRF 209
 Example: DHCP Client and DHCP Server are in Different Tenant VRFs 212
 Example: DHCP Server is in a non-Default, non-VXLAN VRF and DHCP Client is in the Tenant VRF 217
 Additional References for DHCP Relay in a BGP EVPN VXLAN Fabric 221

CHAPTER 7

Configuring VXLAN-Aware Flexible Netflow 223
 Restrictions for VXLAN-Aware Flexible NetFlow 223
 Information About VXLAN-Aware Flexible NetFlow 223
 How to Configure VXLAN-Aware Flexible NetFlow 223
 Configuring a Flow Record 224
 Configuring Flow Exporter 225
 Configuring a Flow Monitor 227
 Configuring Flexible NetFlow on an NVE Interface 227
 Configuration Examples for VXLAN-Aware Flexible NetFlow 228

CHAPTER 8

Configuring BGP EVPN VXLAN QoS 235
 Quality of Service Components 235
 How to Configure BGP EVPN VXLAN QoS 236
 BGP EVPN VXLAN QoS Modes 237
 BGP EVPN VXLAN Traffic without QoS Policy 238
 BGP EVPN VXLAN QoS Marking with an Input Policy 239

BGP EVPN VXLAN QoS Marking with Input and Output Policies 240

CHAPTER 9

Configuring Optimized Layer 2 Overlay Multicast 241

Restrictions for Optimized Layer 2 Overlay Multicast 241

Information About Optimized Layer 2 Overlay Multicast 241

Optimized Layer 2 Multicast in BGP EVPN VXLAN Layer 2 Overlay 242

Optimized Layer 2 Overlay Multicast Interworking in BGP EVPN VXLAN Fabric 243

Optimized Layer 2 Multicast across Layer 2 Overlays 243

Optimized Layer 2 Overlay Multicast Handoff with External Network 244

Replication Types Supported in Underlay Network 245

How to Configure Optimized Layer 2 Overlay Multicast 245

Configure Optimized Layer 2 Multicast in EVPN VXLAN Layer 2 Overlay 245

Enable Optimized Layer 2 Overlay Multicast in BGP EVPN VXLAN Fabric 245

Configure Optimized Layer 2 Overlay Multicast for IPv4 246

Configure Optimized Layer 2 Overlay Multicast for IPv6 247

Configure Optimized Layer 2 Overlay Multicast Interworking 248

Enable Layer 3 Multicast on Centralized Gateway 248

Verifying Optimized Layer 2 Overlay Multicast on VTEP and Centralized Gateway 249

Configuration Examples for Optimized Layer 2 Overlay Multicast 250

Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication 251

Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay 258

Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 266

Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic 273

Example 3: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication 293

Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay 300

Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication 310

Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic 318

Example 5: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 350

Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric 357

Example 6: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6	377
Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric	386

CHAPTER 10**Configuring Tenant Routed Multicast 399**

Restrictions for Tenant Routed Multicast	399
Information about Tenant Routed Multicast	399
TRM in PIM Sparse Mode	402
RP Placement	402
PIM Sparse Mode with Anycast RP	406
PIM Sparse Mode with RP Inside the BGP EVPN VXLAN Fabric	407
PIM Sparse Mode with RP Outside the BGP EVPN VXLAN Fabric	408
TRM in PIM Source Specific Mode	410
Data MDT	411
How to Configure Tenant Routed Multicast	413
Configuring TRM with PIM Sparse Mode	414
Configuring the TRM Multicast Distribution Tree in the VRF	414
Configuring Multicast Routing on the Overlay VRF	415
Configuring Multicast on Switch Virtual Interfaces for Core-facing and Access-facing VLANs	416
Configuring BGP with MVPN Address Family on VTEP	417
Configuring RP for Underlay Network	418
Configuring RP for Overlay Network	419
Configuring TRM with PIM Source Specific Mode	420
Configuring SSM for Overlay Network	420
Verifying Tenant Routed Multicast	421
Troubleshooting Tenant Routed Multicast	422
Configuration Examples for Tenant Routed Multicast	422
Example: Configuring TRM in PIM Sparse Mode with Anycast RP	422
Example: Configuring TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric	461
Example: Configuring TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when the RP is Inside the BGP EVPN VXLAN Fabric	494
Example: Configuring TRM with PIM-SM for IPv4 Multicast Traffic when the RP is Outside the BGP EVPN VXLAN Fabric	533

Example: Configuring TRM with PIM-SM for IPv4 and IPv6 Multicast Traffic when RP is Outside the BGP EVPN VXLAN Fabric 564

CHAPTER 11

Configuring EVPN VXLAN External Connectivity	605
Restrictions for EVPN VXLAN External Connectivity	605
Information About EVPN VXLAN External Connectivity	605
Implementation of Border Nodes for EVPN VXLAN External Connectivity	606
External Connectivity with Layer 3 Networks	608
External Connectivity with Layer 2 Networks	609
Route Leaking between EVPN-enabled VRF and Global Routing Tables	610
How to Configure EVPN VXLAN External Connectivity	610
Enabling Layer 3 External Connectivity with VRF-Lite	611
Configuring the VRF on the Border VTEP Interface that Faces the External Router	611
Enabling Route Leaking between EVPN-enabled VRF and Global Routing Table	612
Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN	614
Configuring BGP on a Border VTEP for External Connectivity with MPLS Layer 3 VPN	615
Enabling EVPN VXLAN Layer 3 TRM Interworking with MVPN Networks	619
Enabling Layer 2 External Connectivity with IEEE 802.1Q Networks	620
Enabling Layer 2 External Connectivity with a VPLS Network Through an Access VFI	621
Defining an Access VFI on a Border VTEP	622
Adding an Access VFI and an EVPN Instance as Members of the VLAN of a Border VTEP	623
Configuring VPLS on a Border VTEP	624
Configuration Examples for EVPN VXLAN External Connectivity	624
Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP	624
Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP	635

CHAPTER 12

Configuring Multi-Homing in a BGP EVPN VXLAN Fabric	647
Restrictions for Multi-Homing in a BGP EVPN VXLAN Fabric	647
Information About Multi-Homing in a BGP EVPN VXLAN Fabric	647
Single-Homing	648
Multi-Homing	649
Single-Active Redundancy Mode	650
Migration Between Single-Homed and Multi-Homed Network Topologies	653
How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric	653

Configuring Dual-Homing with Single Active Redundancy in a BGP EVPN VXLAN Fabric	653
Configuring Redundancy on an Ethernet Segment	653
Associating an Ethernet Segment with an Interface on a VTEP	654
Migrating from a Single-Homed Topology to a Single-Active Dual-Homed Topology	655
Migrating from a Single-Active Dual-Homed Topology to a Single-Homed Topology	657
Configuration Examples for Multi-Homing in a BGP EVPN VXLAN Fabric	658
Example: Configuring Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric	658

CHAPTER 13**Configuring Private VLANs in a BGP EVPN VXLAN Fabric 681**

Restrictions for Private VLANs in a BGP EVPN VXLAN Fabric	681
Information About Private VLANs in a BGP EVPN VXLAN Fabric	681
Primary and Secondary VLANs	682
Private VLAN Ports	682
Extension of Private VLANs in a BGP EVPN VXLAN Fabric	682
Traffic Forwarding for Private VLANs in a BGP EVPN VXLAN Fabric	683
Known Unicast Traffic Forwarding	683
Broadcast, Unknown Unicast, and Multicast Traffic Forwarding	685
Routed Traffic Forwarding	687
How to Configure Private VLANs in a BGP EVPN VXLAN Fabric	687
Configuring the Primary and Secondary VLANs for a Private VLAN	688
Configuring the Port of a Private VLAN	689
Enabling EVPN in a Private VLAN	690
Configuration Examples for Private VLANs in a BGP EVPN VXLAN Fabric	691

CHAPTER 14**Configuring BGP EVPN VXLAN over IPsec 715**

Restrictions for BGP EVPN VXLAN over IPsec	715
Information about BGP EVPN VXLAN over IPsec	715
Workflow to Configure BGP EVPN VXLAN over IPsec	716
Configuration Example for BGP EVPN VXLAN over IPsec	718
Verifying BGP EVPN VXLAN over IPsec	723

CHAPTER 15**Configuring BGP EVPN VXLANv6 731**

Restrictions for BGP EVPN VXLANv6	731
-----------------------------------	-----

Information About BGP EVPN VXLANv6	731
BGP EVPN VXLANv6 Fabric Support	732
EVPN VXLANv6 Layer 2 Overlay	732
EVPN VXLANv6 Overlay with Distributed Anycast Gateway	733
How to Configure BGP EVPN VXLANv6	733
Configure VXLANv6 for BGP EVPN Fabric with Layer 2 Overlay	733
Configure VXLANv6 for BGP EVPN Fabric with Distributed Anycast Gateway	734
Migration from VXLANv4 to VXLANv6	736
Troubleshooting BGP EVPN VXLANv6	740
Configuration Examples for BGP EVPN VXLANv6	740
Example: Configuring BGP EVPN VXLAN with IPv6 Underlay	740
Verifying BGP EVPN VXLAN with IPv6 Underlay Configuration	748
Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay	777
Verifying BGP EVPN VXLAN with Dual Stack Underlay Configuration	787

CHAPTER 16**Configuring EVPN Microsegmentation 821**

Restrictions for EVPN Microsegmentation	821
Information About EVPN Microsegmentation	821
Classification	822
Propagation	822
Enforcement	823
Security Group Access Control Lists	823
How to Configure EVPN Microsegmentation	824
Deployment of EVPN Microsegmentation	824
Leaf Node at Access or Distribution Layer	826
Device Tracking for a Trunk Interface	827
Configuration Example of EVPN Microsegmentation	828

CHAPTER 17**Configuring Additional BGP EVPN Features 829**

Dynamic BGP Peering for EVPN	829
Configuration Example for Dynamic BGP Peering	829
EVPN Route Map Support	830
Configuring Route Maps	830
Matching based on IP prefix	830

- Matching Based on EVPN Route Type 831
- Applying Route Maps 832
- BGP EVPN VRF Auto Route Distinguisher 833
 - Prerequisites for Configuring Auto Route Distinguisher in a BGP EVPN VXLAN Fabric 833
 - Configuring Auto Route Distinguisher in a BGP EVPN VXLAN Fabric 833
- BGP EVPN VRF Auto Route Target 834
 - Prerequisites for Configuring Auto Route Target in a BGP EVPN Fabric 835
 - Configuring Auto Route Target in a BGP EVPN VXLAN Fabric 835

CHAPTER 18 BGP EVPN VXLAN Scalability Guide 837

- Scale and Performance Capabilities for BGP EVPN VXLAN on Cisco Catalyst 9300 Series Switches 837

CHAPTER 19 Troubleshooting BGP EVPN VXLAN 841

- Troubleshooting Scenarios for BGP EVPN VXLAN 841
- Troubleshooting Broadcast, Unknown Unicast, Multicast Traffic Forwarding 842
- Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI 846
 - Verifying the Provisioning of an EVPN VXLAN Layer 2 Overlay Network 847
 - Verifying Intra-Subnet Traffic Movement in an EVPN VXLAN Layer 2 Overlay Network 851
- Troubleshooting Unicast Forwarding Between VTEPs in Different VLANs Through a Layer 3 VNI 858
 - Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network 859
 - Verifying Inter-Subnet Traffic Movement and Symmetric IRB in an EVPN VXLAN Layer 3 Overlay Network 864
- Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network 871
 - Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network 871
 - Verifying Traffic from a VXLAN Fabric to an IP Network Through a Border Leaf Switch Using Route Type 5 871
- Troubleshooting Tenant Routed Multicast 874

PART I Integration with Cisco DNA Service for Bonjour 875

CHAPTER 20 Cisco DNA Service for Bonjour Solution Overview 877

- About the Cisco DNA Service for Bonjour Solution 877
- Solution Components 878
- Supported Platforms 879

Supported Network Design	880
Traditional Wired and Wireless Networks	880
Wired Networks	881
Wireless Networks	883
Cisco SD-Access Wired and Wireless Networks	884
BGP EVPN Networks	886

CHAPTER 21

Configuring Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks	889
Prerequisites for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks	889
Restrictions for Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks	890
Information About Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks	890
Understanding Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks	891
How to Configure Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks	893
Configuring Local Area Bonjour in Unicast Mode over EVPN VXLAN Layer 3 Overlay Networks	893
Configuring Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks	900
Verifying Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks	902
Verifying Local Area Bonjour over an EVPN VXLAN Layer 3 Overlay Network	902
Verifying Wired Service-Peer Configuration	902
Verifying Wired SDG Agent Configuration and Service-Routing Status	904
Verifying Wireless Service-Peer and Service Routing Status	905
Verifying Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks	907
Verifying Cisco DNA Center Configuration and Service-Routing Status	908
Additional References for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks	909

CHAPTER 22

Configuring VRF-Aware Local Area Bonjour Services	911
Prerequisites for VRF-Aware Local Area Bonjour Services	911
Restrictions for VRF-Aware Local Area Bonjour Services	912
Information about VRF-Aware Local Area Bonjour Services	912
Gateway Modes for VRF-Aware Bonjour Services	913
Understanding VRF-Aware Wide Area Bonjour Services	914
Understanding VRF-Aware Service on Multilayered Wired and Wireless Networks	915
How to configure Intra-Virtual Network Proxy Service on Local Area Bonjour Domain	917
How to configure Inter-Virtual Network Proxy Service on Local Area Bonjour Domain	918

Configuring Inter-Virtual Network Location-Filter 919
Verifying VRF-Aware Local Area Bonjour Services 921

PART II

Feature History for BGP EVPN VXLAN 925

CHAPTER 23

Feature History for BGP EVPN VXLAN 927

Feature History for BGP EVPN VXLAN 927



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 Catalyst 9000 Series 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 (bess¹) 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



356449

Overlay Multicast

Overlay multicast is the method by which an 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



356449

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




356448

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

For more information about IRB, see [Information About EVPN VXLAN Integrated Routing and Bridging](#) section.

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 whereas 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 Catalyst 9000 Series 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

- [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](#)

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

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.

See [Example: Configuring BUM Traffic Rate Limiting, on page 71](#) for a sample topology and configuration example.

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 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.

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.
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:	(Optional) Enables default gateway advertisement on the switch. To configure distributed anycast gateway in a VXLAN

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

	Command or Action	Purpose
Step 12	l2vpn evpn instance <i>evpn-instance-number</i> 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.
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.

	Command or Action	Purpose
Step 19	rd <i>rd-value</i> Example: Device(config-evpn-evi) # rd 65000:100	(Optional) Configures a route distinguisher manually.
Step 20	route-target { import export both } <i>rt-value</i> 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 <i>vlan-id</i> Example: Device(config)# vlan configuration 11	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 4	member evpn-instance <i>evpn-instance-id vni l2-vni-number</i> 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 <i>interface-name</i> Example: Device(config)# interface GigabitEthernet1/0/1	Enters interface configuration mode for the specified interface.
Step 4	switchport access vlan <i>vlan-id</i> 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.

	Command or Action	Purpose
Step 3	interface <i>loopback-interface-id</i> Example: Device(config)# interface Loopback0	Enters interface configuration mode for the specified Loopback interface.
Step 4	ip address <i>ipv4-address</i> 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 <i>nve-interface-id</i> 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 <i>loopback-interface-id</i> Example: Device(config-if)# source-interface loopback0	Sets the IP address of the specified loopback interface as the source IP address.

	Command or Action	Purpose
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 Catalyst 9000 Series 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.

	Command or Action	Purpose
Step 3	router bgp <i>autonomous-system-number</i> 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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 5	bgp update-delay <i>time-period</i> 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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 8	neighbor <i>ip-address remote-as number</i> 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 { <i>ip-address</i> <i>group-name</i> } update-source <i>interface</i> 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 <i>l2vpn evpn</i> Example: Device (config-router) # address-family l2vpn evpn	Specifies the L2VPN address family and enters address family configuration mode.

	Command or Action	Purpose
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.
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.

Command	Purpose
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 <i>evpn-instance-id</i> 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 <i>bridge-domain-number</i> detail	Displays detailed information for a Layer 2 forwarding information base bridge domain.
show l2fib bridge-domain <i>bridge-domain-number</i> 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 <i>vni-id</i> 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 <i>vlan-id</i>	Displays MAC addresses for a VLAN.
show platform software fed switch active matm macTable vlan <i>vlan-id</i>	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.

Command	Purpose
<code>show ip mroute</code>	Displays multicast routing table information.
<code>show ip bgp l2vpn evpn detail l2vpn-evpn-route</code>	Displays detailed information about a specific route.
<code>show ip bgp l2vpn evpn detail [mac-address ip-address]</code>	Display routes containing an IP address only or both MAC address and IP address.
<code>show ip bgp l2vpn evpn route-type 2 ethernet-tag {mac-address }</code>	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



357016

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
<pre> Leaf-01# show running-config hostname Leaf-01 ! 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 GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface TenGigabitEthernet1/1/1 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 ! </pre>	<pre> Leaf-02# show running-config hostname Leaf-02 ! 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 GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface TenGigabitEthernet1/1/1 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 ! </pre>

VTEP 1	VTEP 2
<pre> 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 Leaf-01# </pre>	<pre> 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 Leaf-02# </pre>

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:

```

Leaf-01# 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

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 172.16.255.1, 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)
    
```

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

```

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

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 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 ?

Leaf-01#

```

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

```

Leaf-01# 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

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

(*, 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:
  TenGigabitEthernet1/1/1, 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:
  TenGigabitEthernet1/1/1, 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:
  TenGigabitEthernet1/1/1, 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: TenGigabitEthernet1/1/1, RPF nbr 172.16.12.2
Outgoing interface list:
  Tunnel0, Forward/Sparse-Dense, 00:00:03/00:02:56

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
Tunnel2 Flags: A
TenGigabitEthernet1/1/1 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

```

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

```

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
TenGigabitEthernet1/1/1 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
TenGigabitEthernet1/1/1 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
TenGigabitEthernet1/1/1 Flags: NS

```

```
Leaf-01#
```

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:

```
Leaf-02# 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
```

```
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 172.16.255.2, 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
```

```
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 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 ?

```

Leaf-02#

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

```

Leaf-02# 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

```

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 - encaps-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: TenGigabitEthernet1/1/1, RPF nbr 172.16.12.1
Outgoing interface list:
Loopback0, Forward/Sparse, 00:43:49/00:02:09

```

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

```
(*, 225.0.0.101), 00:43:49/stopped, RP 172.16.255.1, flags: SJCFx
  Incoming interface: TenGigabitEthernet1/1/1, 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: TenGigabitEthernet1/1/1, 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:
    TenGigabitEthernet1/1/1, Forward/Sparse, 00:00:20/00:03:09
```

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
  TenGigabitEthernet1/1/1 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
  TenGigabitEthernet1/1/1 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
  TenGigabitEthernet1/1/1 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
  TenGigabitEthernet1/1/1 Flags: F NS
```



```

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

```


Leaf-02#

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



957017

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
<pre> Leaf-01# show running-config hostname Leaf-01 ! 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 GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface TenGigabitEthernet1/1/1 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 ! </pre>	<pre> Leaf-02# show running-config hostname Leaf-02 ! 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 GigabitEthernet1/0/10 switchport access vlan 101 switchport mode access spanning-tree portfast ! interface TenGigabitEthernet1/1/1 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 ! </pre>

VTEP 1	VTEP 2
<pre> 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 Leaf-01# </pre>	<pre> 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 Leaf-02# </pre>

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:

```

Leaf-01# 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

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 172.16.255.1, 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
    
```

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

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

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 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 ?

Leaf-01#
```

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

```
Leaf-01# 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

Leaf-01#
```

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

```
Leaf-01# 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

```

Leaf-01#

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:

```

Leaf-02# 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

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 172.16.255.2, 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

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 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

```

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

```

      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 ?

Leaf-02#

```

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

```

Leaf-02# 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

```

```
Leaf-02#
```

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

```

Leaf-02# 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

```


Leaf-02#

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
<pre> Leaf-01# show running-config hostname Leaf-01 ! 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 ! </pre>	<pre> Leaf-02# show running-config hostname Leaf-02 ! 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 ! </pre>

VTEP 1	VTEP 2
<pre>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 Leaf-01#</pre>	<pre>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 Leaf-02#</pre>

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

Spine Switch 1	Spine Switch 2
<pre> 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 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 ! </pre>	<pre> 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 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 ! </pre>

Spine Switch 1	Spine Switch 2
<pre> 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 Spine-01# </pre>	<pre> 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 Spine-02# </pre>

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:

```

Leaf-01# 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

Leaf-01#

```

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

```

Leaf-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
       & - replicated local route overrides by connected

```

Example: Configuring Layer 2 VNI with Spine Multicast Replication

```

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

```

```
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 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 rinfo 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

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 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 ?
Leaf-01#

```

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

```

Leaf-01# 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
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

(*, 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:

```

Example: Configuring Layer 2 VNI with Spine Multicast Replication

```

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

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
  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

```

```
Leaf-01#
```

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:

```

Leaf-02# 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

Leaf-02#

```

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

```

Leaf-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
       & - 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

```

Example: Configuring Layer 2 VNI with Spine Multicast Replication

```

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

```

```
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 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

```

```
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 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 ?

```

```
Leaf-02#
```


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

```
Leaf-02# 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

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

(*, 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
    GigabitEthernet1/0/1, 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

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
```

Example: Configuring Layer 2 VNI with Spine Multicast Replication

```

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
  GigabitEthernet1/0/1 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
  GigabitEthernet1/0/2 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:  146/0/167/0, Other: 0/0/0
  GigabitEthernet1/0/2 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
  GigabitEthernet1/0/2 Flags: F NS
    Pkts: 0/0/0   Rate: 0 pps

```

Leaf-02#

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:

```

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, 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

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 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

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 35, local router ID is 172.16.255.1

```

Example: Configuring Layer 2 VNI with Spine Multicast Replication

```

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 ?

```

```
Spine-01#
```

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

```

Spine-01# 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      ?
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), 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: GigabitEthernet1/0/3, 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: GigabitEthernet1/0/2, RPF nbr 172.16.13.3
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: 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
Tunnell Flags: A
GigabitEthernet1/0/3 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
Tunnell 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
Tunnell 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
Tunnell Flags: A
GigabitEthernet1/0/3 Flags: NS

Spine-01#

```

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:

```
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, 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

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 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

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 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	?

Spine-02#

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

```
Spine-02# 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      ?
```

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,
```

Example: Configuring Layer 2 VNI with Spine Multicast Replication

```

    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: GigabitEthernet1/0/3, RPF nbr 172.16.24.4
Outgoing interface list:
  GigabitEthernet1/0/2, 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

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
Tunnell Flags: A
GigabitEthernet1/0/3 Flags: F IC NS
  Pkts: 0/0/0   Rate: 0 pps
GigabitEthernet1/0/2 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
Tunnell Flags: A
GigabitEthernet1/0/3 Flags: F NS
Pkts: 0/0/2    Rate: 0 pps
GigabitEthernet1/0/2 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
GigabitEthernet1/0/3 Flags: A
GigabitEthernet1/0/2 Flags: F NS
Pkts: 0/0/0    Rate: 0 pps
    
```


Spine-02#

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



357015

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
<pre>Leaf-01# show running-config hostname Leaf-01 ! 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 !</pre>	<pre>Leaf-02# show running-config hostname Leaf-02 ! 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 !</pre>

VTEP 1	VTEP 2
<pre>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 Leaf-01#</pre>	<pre>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 Leaf-02#</pre>

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

Spine Switch 1	Spine Switch 2
<pre> 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 ! </pre>	<pre> 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 ! </pre>
<pre> 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# </pre>	<pre> 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# </pre>

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:

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

Leaf-01#
```

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

```
Leaf-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
       & - 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
```

Example: Configuring Layer 2 VNI with Spine Ingress Replication

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

```
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 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

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 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 ?

Leaf-01#

```

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

```

Leaf-01# 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

Leaf-01#

```

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

```

Leaf-01# 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

Leaf-01#

```

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:

```

Leaf-02# 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

Leaf-02#

```

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

Example: Configuring Layer 2 VNI with Spine Ingress Replication

```
Leaf-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
       & - 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
```

```
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 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 rinfo 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

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 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 ?

Leaf-02#
```

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

```
Leaf-02# 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

Leaf-02#
```

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

```
Leaf-02# 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
```

Example: Configuring Layer 2 VNI with Spine Ingress Replication

```

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

```

```
Leaf-02#
```

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 rinfo 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.

Example: Configuring BUM Traffic Rate Limiting

This example shows how to configure and verify BUM traffic rate limiting in a BGP EVPN VXLAN fabric using the following topology:



356465

The topology shows an EVPN VXLAN network with 2 VTEPs (VTEP 1 and VTEP 2) connected to perform bridging.

Configuring BUM Traffic Rate Limiting on a VTEP

The following example provides a sample configuration for BUM traffic rate limiting on VTEP 1:

```
Leaf-01# configure terminal
Leaf-01(config)# class-map match-all CL2Miss
Leaf-01(config-cmap)# match 12 dst-mac miss
Leaf-01(config-cmap)# exit
Leaf-01(config)# policy-map PL2Miss
Leaf-01(config-pmap)# class CL2Miss
Leaf-01(config-pmap-c)# police 100000
Leaf-01(config-pmap-c)# exit
Leaf-01(config)# interface nve1
Leaf-01(config-if)# service-policy output PL2Miss
Leaf-01(config-if)# exit
Leaf-01(config)# end
Leaf-01#
```

Verifying BUM Traffic Rate Limiting on a VTEP

The following example shows how to check the aggregated policy map and rate statistics on VTEP 1:

```
Leaf-01# show policy-map int nve1

nve1
  Service-policy output: PL2Miss
  Class-map: sam1 (match-all)
    0 packets
  Match: 12 dst-mac miss
  police:
    cir 100000 bps, bc 3125 bytes
    conformed 221238 bytes; actions:
      transmit
    exceeded 2647233234 bytes; actions:
      drop
    conformed 7000 bps, exceeded 69060000 bps
  Class-map: class-default (match-any)
    10022668 packets
  Match: any

Leaf-01#
```

The following example shows how to validate the member VNI policy under an NVE on VTEP 1:

```
Leaf-01# show platform software fed switch active qos policy target brief | begin PL2Miss
TCG summary for policy: PL2Miss
Loc Interface IIF-ID Dir tccg Child #m/p/q State:(cfg,opr)
-----
L:255 nve1.VNI10000 0x00000000420012 OUT 2 0 0/1/0 VALID,SET_INHW
0x7f605dc9b258
L:255 nve1 0x000000000000bb OUT 2 0 0/1/0 VALID,INIT
0x7f605dc9c2f8

Leaf-01#
```

The following example shows how to validate the individual statistics on VTEP 1:

```
Leaf-01# show platform software fed switch active qos policer all_instances trail
All policer instances: With trail
*****
List of AAL QoS Policer Instances on Targets
AAL Info:
=====
Handle : 0x4
Target : 0xdf0001b7(iif_id : 0x420012)
Asic num : 0x0
Policer Type : Aggregate
le id : 0x5db76438
le Type : PORT
Ingress Block: 0x0
Egress Block : 0x25
Policer HW info:
  Ingress:(Total : 0)
    Policer Policer Policer
    Number Type offset
    -----
  Egress:(Total : 1)
    Policer Policer Policer
    Number Type offset
    -----
          0 1R2C 0
RAL handle : 4294967295
RAL Info:(Base:Double)
=====
AFD handles : Ingress - Not allocated Egress - 0

AFD QIM Info:
=====
Policer Block Handle : 0
ASIC Num : 0(Physical:0, Core 0)
LE ID : 278
LE Type : 1
Policer Base : 126976
Size : 1
Start Index : 0
End Index : 0
Ingress Offset : 1
Ingress Offsets : 1R2C:0(Total:0), 1R3C:0(Total:0), 2R3C:0(Total:0)
Egress Offsets : 1R2C:0(Total:1), 1R3C(Total:0):0, 2R3C:0(Total:0)

Policer|Policer|Rate |Exceed Rate |Burst Size |Exceed Burst
Size |Drop or |Exceed Drop|Mark Tbl |Class |Color|Offset |Type |(bps) [RegVal]
| (bps) [RegVal] | (Bytes) [RegVal] | (Bytes) [RegVal]
|Markdown|orMarkdown|Exceed/Violate|Default|Aware|
=====
0|Out1R2C|100057 [0x2f3b] |29 [0x0000] |3136 [49] |0 [0]
```

Example: Configuring BUM Traffic Rate Limiting

```

|DROP |N/A |0x0/0x0 |No |No |
-----|-----|-----|-----|-----|
Policer|DMA Stats In (Bytes) |DMA Stats Out (Bytes) |DMA Stats In (Frames)
|DMA Stats Out (Frames) |Offset |Green/Yellow |DMA Stats In (Frames)
|Green/Yellow |Green/Yellow/Red |Green/Yellow/Red
-----|-----|-----|-----|-----|
0| 2647454472/ | 0| 221238/ 2647233234/ | 0| 25955436/
0| 2169/ 25953267/ | 0|
-----|-----|-----|-----|-----|
***** END *****

Leaf-01#

```




CHAPTER 3

Configuring EVPN VXLAN Layer 3 Overlay Network


- [Information About EVPN VXLAN Layer 3 Overlay Network](#), on page 75
- [How to Configure EVPN VXLAN Layer 3 Overlay Network](#), on page 76
- [Configuration Examples for EVPN VXLAN Layer 3 Overlay Network](#), on page 86
- [Verifying EVPN VXLAN Layer 3 Overlay Network](#), on page 93

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 <i>vrf-name</i> Example: Device (config)# vrf definition Green	Enters the VRF configuration mode for the specified VRF instance.

	Command or Action	Purpose
Step 4	rd <i>vpn-route-distinguisher</i> 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 } <i>route-target-ext-community</i> 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 } <i>route-target-ext-community</i> 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 } <i>route-target-ext-community</i> 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} <i>route-target-ext-community stitching</i> 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: Device(config-vlan)# end	Returns to privileged EXEC mode.

	Command or Action	Purpose
--	-------------------	---------

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 <i>interface-name</i> Example: Device(config)# interface GigabitEthernet1/0/1	Enters interface configuration mode for the specified interface.
Step 4	switchport access vlan <i>vlan-id</i> 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.
Step 2	configure terminal Example:	Enters global configuration mode.

	Command or Action	Purpose
	Device# configure terminal	
Step 3	interface vlan <i>vlan-id</i> Example: Device(config)# interface vlan 11	Enters interface configuration mode for the specified VLAN.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding Green	Configures the SVI for the VLAN.
Step 5	ip unnumbered <i>Loopback-interface</i> 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.
Step 3	interface vlan <i>vlan-id</i> Example: Device(config)# interface vlan 40	Enters interface configuration mode for the specified VLAN.

	Command or Action	Purpose
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding Green	Configures the SVI for the VLAN.
Step 5	ip address <i>ip-address</i> 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 <i>mac-address-value</i> 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 <i>loopback-interface-id</i> Example: Device(config)# interface Loopback0	Enters interface configuration mode for the specified Loopback interface.
Step 4	ip address <i>ipv4-address</i> Example:	Configures the IP address for the Loopback interface.

	Command or Action	Purpose
	Device (config-if) # <code>ip address 10.12.11.11 255.255.255.255</code>	
Step 5	ip pim sparse mode Example: Device (config-if) # <code>ip pim sparse mode</code>	(Optional) Enables Protocol Independent Multicast (PIM) sparse mode 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 6	end Example: Device (config-vlan) # <code>end</code>	Returns to privileged EXEC mode.

Configuring the NVE Interface on a VTEP

To add a Layer 3 VNI member to the NVE interface on a VTEP, perform the following steps:

Procedure

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

	Command or Action	Purpose
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 Catalyst 9000 Series switches.
Step 7	member vni vni-id vrf vrf-name Example: Device(config-if)# member vni 5000 vrf Green	Associates the Layer 3 VNI id with the NVE interface. Note The Layer 3 VNI id must match with the VNI id configured in the core VLAN on the VTEP.
Step 8	end Example: Device(config-if)# end	Returns to privileged EXEC mode.

Configuring BGP with IPv4 or IPv6 or Both Address Families on VTEP

To configure BGP on a VTEP with IPv4 or IPv6 or both address families and a 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.


	Command or Action	Purpose
Step 4	bgp log-neighbor-changes Example: <pre>Device(config-router)# bgp log-neighbor-changes</pre>	(Optional) Enables the generation of logging messages when the status of a BGP neighbor changes. For more information, see <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 5	bgp update-delay time-period Example: <pre>Device(config-router)# bgp update-delay 1</pre>	(Optional) Sets the maximum initial delay period before sending the first update. For more information, see <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 6	bgp graceful-restart Example: <pre>Device(config-router)# bgp graceful-restart</pre>	(Optional) Enables the BGP graceful restart capability for all BGP neighbors. For more information, see <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 7	no bgp default ipv4-unicast Example: <pre>Device(config-router)# no bgp default ipv4-unicast</pre>	(Optional) Disables default IPv4 unicast address family for BGP peering session establishment. For more information, see <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 8	neighbor ip-address remote-as number Example: <pre>Device(config-router)# neighbor 10.11.11.11 remote-as 1</pre>	Defines multiprotocol-BGP neighbors. Under each neighbor, define the 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: <pre>Device(config-router)# neighbor 10.11.11.11 update-source Loopback0</pre>	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: <pre>Device(config-router)# address-family l2vpn evpn</pre>	Specifies the L2VPN address family and enters address family configuration mode.
Step 11	neighbor ip-address activate Example: <pre>Device(config-router-af)# neighbor 10.11.11.11 activate</pre>	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:	Specifies the communities attribute sent to a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.

	Command or Action	Purpose
	Device(config-router-af)# neighbor 10.11.11.11 send-community both	
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 Example: Device(config-router-af)# advertise l2vpn evpn	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric.
Step 16	redistribute connected Example: Device(config-router-af)# redistribute connected	(Optional) Redistributes connected routes to BGP.
Step 17	redistribute static Example: Device(config-router-af)# redistribute static	(Optional) 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 Example: Device(config-router-af)# advertise l2vpn evpn	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric.
Step 21	redistribute connected Example: Device(config-router-af)# redistribute connected	(Optional) Redistributes connected routes to BGP.

	Command or Action	Purpose
Step 22	redistribute static Example: Device (config-router-af) # redistribute static	(Optional) 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.



356465



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
<pre> 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 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 GigabitEthernet1/0/29 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 </pre>	<pre> 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 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 GigabitEthernet1/0/30 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 </pre>

VTEP 1	VTEP 2
<pre> 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 </pre>	<pre> 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 </pre>

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

- [#unique_59 unique_59_Connect_42_section_zll_qxs_nkb](#)
- [#unique_59 unique_59_Connect_42_section_zwz_pxs_nkb](#)
- [#unique_59 unique_59_Connect_42_section_y3n_pxs_nkb](#)
- [#unique_59 unique_59_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:

```
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 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:

```
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/-/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** command on VTEP 1:

```
VTEP1# 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:


```

VTEP2# 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:

```

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, 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:

```

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, 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.

```

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      0      0 Vlan200
200   380e.4d9b.6a4a                       0x1000001   0   0      64 0x7f5d85117598      0x7f5d85110f78
      0x7f5d851b9648                       0x0      0      0      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.

```
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  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 4

Configuring EVPN VXLAN Integrated Routing and Bridging

- [Restrictions for EVPN VXLAN Integrated Routing and Bridging, on page 95](#)
- [Information About EVPN VXLAN Integrated Routing and Bridging, on page 95](#)
- [How to Configure EVPN VXLAN Integrated Routing and Bridging, on page 101](#)
- [Verifying EVPN VXLAN Integrated Routing and Bridging, on page 109](#)
- [Configuration Examples for EVPN VXLAN Integrated Routing and Bridging, on page 110](#)

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.
- 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 only supported starting from Cisco IOS XE Amsterdam 17.3.2a release.

Information About EVPN VXLAN Integrated Routing and Bridging

EVPN VXLAN integrated routing and bridging (IRB) allows the VTEPs or leaf switches in an EVPN VXLAN network to perform both bridging and routing. IRB allows the VTEPs to forward both Layer 2 or bridged and

Layer 3 or routed traffic. A VTEP performs bridging when it forwards traffic to the same subnet. Similarly, a VTEP performs routing when it forwards traffic to a different subnet. The VTEPs in the network forward traffic to each other through the VXLAN gateways. BGP EVPN VXLAN implements IRB in two ways:

- Asymmetric IRB
- Symmetric IRB

Asymmetric IRB

In asymmetric IRB, the ingress VTEP performs both bridging and routing whereas the egress VTEP performs only bridging. A packet first moves through a MAC VRF followed by an IP VRF on the network visualization endpoint (NVE) of the ingress VTEP. It then moves only through a MAC VRF on the NVE of the egress VTEP. The NVE of the ingress VTEP manages all the packet processing associated with intersubnet forwarding semantics.


The return traffic during asymmetric IRB goes through a different virtual network instance (VNI) compared to the source traffic. Asymmetric IRB needs the source and destination VNIs to be associated with both the ingress and egress VTEPs.

Symmetric IRB

In symmetric IRB, both the ingress and egress VTEPs perform both bridging and routing. A packet first moves through a MAC VRF followed by an IP VRF on the NVE of the ingress VTEP. It then moves through an IP VRF followed by a MAC VRF on the NVE of the egress VTEP. The NVEs of ingress and egress VTEPs equally share all the packet processing associated with intersubnet forwarding semantics.

In symmetric IRB, you are required to define only the VNIs of locally attached endpoints on the ingress and egress VTEPs. Symmetric IRB offers better scalability in terms of the number of VNIs that a BGP EVPN VXLAN fabric supports.

The following figure shows the implementation of symmetric IRB and the movement of traffic in an EVPN VXLAN network:




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 11: Distributed Gateway Topology



When multiple VETPs 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.

In the [Figure 11: Distributed Gateway Topology](#) image, to enable distributed anycast gateway in subnet 1, configure the same MAC address on leaf switch 1 and leaf switch 2. This ensures that the ARP entries of gateway MAC and IP addresses on host device 1 match with the MAC and IP addresses of both leaf switch 1 and leaf switch 2.

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.

In the [Figure 11: Distributed Gateway Topology](#) image, consider the scenario where MAC aliasing is enabled in subnet 1. Leaf switch 1 and leaf switch 2 advertise their MAC addresses to each other as gateway MAC addresses. This allows leaf switch 2 to recognize the MAC address in the ARP entry of host device 1 as a gateway MAC address. It allows host device 1 to send traffic outside of subnet 1 even though its VLAN MAC address does not match with the ARP entry.

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:

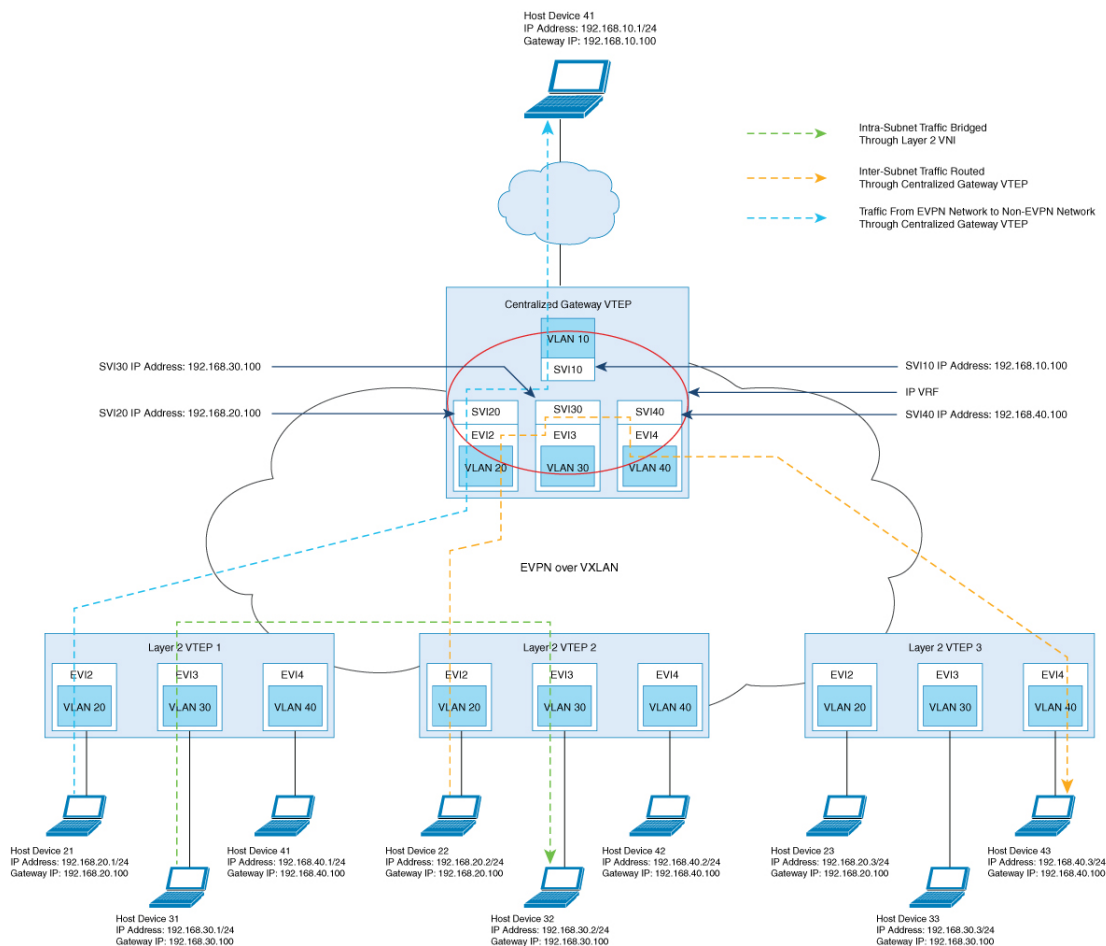
- You require a boundary between Layer 2 and Layer 3 segments at the border of the BGP EVPN VXLAN fabric.
- The inter-subnet traffic is subjected to a firewall inspection or any policy on a centralized plane.



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 12: 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. See [RFC4732](#) for more details about the 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. For detailed steps, see [Configuring Layer 2 VPN EVPN on a VTEP, on page 101](#) and [Configuring the NVE Interface on a VTEP, on page 104](#) sections.



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:

Configuring Layer 2 VPN EVPN on a VTEP

See [Configuring Layer 2 VPN EVPN on a VTEP, on page 17](#) for detailed steps.

Configuring IP VRF on VTEP

See [Configuring an IP VRF on a VTEP, on page 76](#) for detailed steps.

Configuring the Loopback Interface on a VTEP

See [Configuring the Loopback Interface on a VTEP, on page 81](#) for detailed steps.

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 <i>vlan-id</i> Example: Device(config)# vlan configuration 201	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 4	member evpn-instance <i>evpn-instance-id vni l2-vni-number</i> Example:	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 2 VNI.

	Command or Action	Purpose
	<code>Device (config-vlan) # member evpn-instance 1 vni 6000</code>	
Step 5	exit Example: <code>Device (config-vlan) # exit</code>	Returns to global configuration mode.
Step 6	vlan configuration <i>vlan-id</i> Example: <code>Device (config) # vlan configuration 202</code>	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 7	member evpn-instance <i>evpn-instance-id</i> vni <i>l2-vni-number</i> Example: <code>Device (config-vlan) # member evpn-instance 2 vni 7000</code>	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 2 VNI.
Step 8	exit Example: <code>Device (config-vlan) # exit</code>	Returns to global configuration mode.
Step 9	vlan configuration <i>vlan-id</i> Example: <code>Device (config) # vlan configuration 200</code>	Enters VLAN feature configuration mode for the specified VLAN interface.
Step 10	member vni <i>l3-vni-number</i> Example: <code>Device (config-vlan) # member vni 5000</code>	Adds EVPN instance as a member of the VLAN configuration. The VNI here is used as a Layer 3 VNI.
Step 11	exit Example: <code>Device (config-vlan) # exit</code>	Returns to global configuration mode.
Step 12	end Example: <code>Device (config-vlan) # end</code>	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	<code>enable</code>	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 <i>vlan-id</i> Example: Device(config)# interface vlan 200	Enters interface configuration mode for the specified VLAN.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding Green	Configures the SVI for the VLAN.
Step 5	ip unnumbered <i>Loopback-interface</i> 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 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.

	Command or Action	Purpose
Step 3	interface vlan <i>vlan-id</i> Example: Device(config)# interface vlan 202	Enters interface configuration mode for the specified VLAN.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding Green	Configures the SVI for the VLAN.
Step 5	ip address <i>gateway-ip-address</i> 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 <i>mac-address-value</i> 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 <i>nve-interface-id</i> Example: Device(config)# interface nve1	Defines the interface to be configured as a trunk, and enters interface configuration mode.

	Command or Action	Purpose
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 Catalyst 9000 Series 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. 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 <i>autonomous-system-number</i> 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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 5	bgp update-delay <i>time-period</i> Example: Device (config-router)# bgp update-delay 1	(Optional) Sets the maximum initial delay period before sending the first update. For more information, see <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
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 <i>Configuring BGP</i> module of the <i>IP Routing Configuration Guide</i> .
Step 8	neighbor <i>ip-address</i> remote-as <i>number</i> 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 {<i>ip-address</i> <i>group-name</i>} update-source <i>interface</i> Example:	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.

	Command or Action	Purpose
	Device(config-router)# neighbor 10.11.11.11 update-source Loopback0	
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: Device(config-router)# address-family ipv4 vrf green	Specifies the IPv4 address family and enters address family configuration mode.
Step 15	advertise l2vpn evpn Example: Device(config-router-af)# advertise l2vpn evpn	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric.
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:	Exits address family configuration mode and returns to router configuration mode.

	Command or Action	Purpose
	Device (config-router-af) # exit-address-family	
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 Example: Device (config-router-af) # advertise l2vpn evpn	Advertises Layer 2 VPN EVPN routes within a tenant VRF in an EVPN VXLAN fabric.
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.

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 101](#)
 - [Configuring the Loopback Interface on a VTEP, on page 101](#)
 - [Configuring Core-facing and Access-facing VLANs on a VTEP, on page 101](#)
 - [Configuring Switch Virtual Interface for the Core-facing VLAN on a VTEP, on page 102](#)

- [Configuring Switch Virtual Interface for the Access-facing VLANs on a VTEP, on page 103](#)
- [Configuring the NVE Interface on a VTEP, on page 104](#)
- [Configuring BGP with EVPN and VRF Address Families on a VTEP, on page 105](#)

**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 93](#) 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.

Command	Purpose
<code>show mac address-table</code>	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 13: 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
--------	--------

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

VTEP 1	VTEP 2
<pre> 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 ! </pre>	<pre> 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 ! </pre>

VTEP 1	VTEP 2
<pre> ! interface GigabitEthernet1/0/29 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 description vni 6000 default-gateway vrf forwarding green ip address 192.168.1.201 255.255.255.0 ipv6 address 2001:DB8:201::201/64 ipv6 enable ! interface Vlan202 description vni 7000 default-gateway 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 6000 mcast-group 232.1.1.1 member vni 5000 vrf green member vni 7000 mcast-group 232.1.1.1 ! 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 </pre>	<pre> ! interface GigabitEthernet1/0/30 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 Vlan201 description vni 6000 default-gateway vrf forwarding green ip address 192.168.1.201 255.255.255.0 ipv6 address 2001:DB8:201::201/64 ipv6 enable ! interface Vlan202 description vni 7000 default-gateway 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 6000 mcast-group 232.1.1.1 member vni 7000 mcast-group 232.1.1.1 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 </pre>

VTEP 1	VTEP 2
<pre> 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 </pre>	<pre> 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 </pre>

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

- [show nve peers, on page 114](#)
- [show l2vpn evpn peers vxlan, on page 115](#)
- [show l2vpn evpn evi evpn-instance detail, on page 115](#)
- [show l2vpn evpn default-gateway, on page 116](#)
- [show bgp l2vpn evpn all, on page 117](#)
- [show ip route vrf green, on page 120](#)
- [show platform software fed switch active matm mactable vlan, on page 121](#)

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 10.1.1.10      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# show 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 V1201
Y  192.168.1.201                1    201   380e.4d9b.6a48 10.2.2.20
Y  2001:DB8:201::201            1    201   a0f8.4910.bccc V1201
Y  2001:DB8:201::201            1    201   380e.4d9b.6a48 10.2.2.20
Y  192.168.2.202                2    202   a0f8.4910.bcc2 V1202
Y  192.168.2.202                2    202   380e.4d9b.6a42 10.2.2.20
Y  2001:DB8:202::202            2    202   a0f8.4910.bcc2 V1202
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	V1201
Y	192.168.1.201	1	201	a0f8.4910.bccc	10.1.1.10
Y	2001:DB8:201::201	1	201	380e.4d9b.6a48	V1201
Y	2001:DB8:201::201	1	201	a0f8.4910.bccc	10.1.1.10
Y	192.168.2.202	2	202	380e.4d9b.6a42	V1202
Y	192.168.2.202	2	202	a0f8.4910.bccc	10.1.1.10
Y	2001:DB8:202::202	2	202	380e.4d9b.6a42	V1202
Y	2001:DB8:202::202	2	202	a0f8.4910.bccc	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
```

Route	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					

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

```

::: 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 ?

```

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

```

*>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

```

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

```

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 riHandle	Type diHandle	Seq#	EC_Bi	Flags	machandle *a_time *e_time ports	siHandle
201	00aa.00bb.00cc 0x0	0x8002 0x0	0	42949	64	0x7f5d85007b88 0 0 Vlan201	0x7f5d852b6d28
201	0059.dc50.ae01 0x0	0x1 0x7f5d8517eae8	9	0	0	0x7f5d852abaf8 300 9 GigabitEthernet1/0/1	0x7f5d85035248
201	a0f8.4910.bccc 0x0	0x8002 0x5234	0	19880	64	0x7f5d852ad618 0 9 Vlan201	0x7f5d852b6d28
201	0059.dc50.ae4c 0x0	0x1 0x7f5d8517eae8	16	0	0	0x7f5d855b3ff8 300 95 GigabitEthernet1/0/1	0x7f5d855a2858
201	380e.4d9b.6a48 0x0	0x8002 0x5234	0	0	64	0x7f5d84fbf948 0 95 Vlan201	0x7f5d852b6d28
201	0018.736c.56c3 0x7f5d855c6098	0x1000001 0x0	0	0	64	0x7f5d855c8268 0 95 RLOC 10.2.2.20 adj_id 36	0x7f5d852368b8

Total Mac number of addresses:: 6

VTEP1# show platform software fed switch active matm macTable vlan 202

VLAN	MAC riHandle	Type diHandle	Seq#	EC_Bi	Flags	machandle *a_time *e_time ports	siHandle
202	a0f8.4910.bcc2 0x0	0x8002 0x0	0	19880	64	0x7f5d8503d288 0 0 Vlan202	0x7f5d852b6d28
202	0059.dc50.ae01 0x0	0x1 0x7f5d8517eae8	10	0	0	0x7f5d852ac8b8 300 15 GigabitEthernet1/0/1	0x7f5d852ac668
202	0018.736c.5681 0x7f5d8518dea8	0x1000001 0x0	0	0	64	0x7f5d855ba7a8 0 15 RLOC 10.2.2.20 adj_id 125	0x7f5d855b0c58
202	0059.dc50.ae42 0x0	0x1 0x7f5d8517eae8	17	0	0	0x7f5d8518e848 300 225 GigabitEthernet1/0/1	0x7f5d855a5258
202	380e.4d9b.6a42 0x0	0x8002 0x5234	0	0	64	0x7f5d855a59a8 0 225 Vlan202	0x7f5d852b6d28
202	0018.736c.56c2 0x7f5d8518dea8	0x1000001 0x0	0	0	64	0x7f5d8523d2b8 0 225 RLOC 10.2.2.20 adj_id 125	0x7f5d855b0c58

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
```

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

```

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

202    a0f8.4910.bcc2    0x8002      0    0    64    0x7fa88594ddb8    0x7fa885574e38
      0x0          0x5174      0          0          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 14: 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:

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

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
<pre> 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 </pre>	<pre> 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 ! </pre>	<pre> 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 </pre>

Leaf Switch 1	Centralized Gateway Switch	Leaf Switch 2
<pre> ! 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 172.16.255.255 ! end ! Leaf_Switch1# </pre>	<pre> 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 GigabitEthernet1/0/10 switchport mode trunk ! interface Vlan101 description centralized gateway vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan102 description centralized gateway vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 172.16.255.5 ! 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 </pre>	<pre> ! 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 172.16.255.255 ! end ! Leaf_Switch2# </pre>

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

```
en#  
lc#
```


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

```

en S
lc 6

sw
337

S
p i
rip
p i
ip
1
aa
0
!
en
n
p i
sw
137

S
p i
rip
p i
ip
h
p i
ip
1
aa
0
!
en
n

```



```

enps
lcw
  ip
  pi
  scd
  vz
  ip
  mip
  pi
  fso
  km
  ip
  fso
  1
  aar
  0
  !
  chl
  on
  ip
  scd
  vz
  ip
  mip
  pi
  fso
  km
  ip
  fso
  1
  
```

```

en S
ict 6
ata
  0
  !
  xmr
  ipso
  1
  clear
  !
  !
  !
  xmr
  pgb
  !
  !
  pgb
  clear
  !
  !
  pgb
  con
  pgb
  tlat
  !
  !
  327
  clear
  !
  !
  327
  comp
  !
  !
  427
  clear
  !
  !
  427
  comp
  !

```

```

enps
lcw
mgm
527
smr
t5
mgm
527
cmh
t5
!
yca
t5
!

```

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

```
en#  
lc#
```

```

enps
let
!
vfa
ns21
re
rgn
337
dn
rgn
337
vwb
lb
rgn
337
dn
rgn
437
dn
rgn
437
vwb
lb
rgn
437
dn
rgn
537
dn
rgn
537
vwb
lb
rgn
537
dn
rgn
!
pi
  
```

```

enS
lc 6
rip
s
!
!
!
!
!

```

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: 0x40000000000003, 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  Gi1/0/10:101
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
    
```

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

```

-----
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: 0x40000000000003, 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	Gi1/0/10:101
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	Gi1/0/10:102
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 V1101
Y  10.1.102.1                   102  102   7c21.0dbd.274d V1102

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:
```

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

```

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   V1101:0
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   V1102:0
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):          V1101:0

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):          V1102:0

```



CHAPTER 5

Configuring Spine Switches in a BGP EVPN VXLAN Fabric

- [Information About Spine Switches in a BGP EVPN VXLAN Fabric, on page 141](#)
- [Configuration Examples for Spine Switches in a BGP EVPN VXLAN Network, on page 142](#)

Information About Spine Switches in a BGP EVPN VXLAN Fabric

Spine switches in a BGP EVPN VXLAN fabric act as the connecting nodes between all the leaf switches or VTEPs. They form the backbone of the EVPN VXLAN network and forward traffic between the leaf switches. Each leaf switch is connected to each spine switch in the network. Spine switches enable redundancy within the network and provide multiple paths for VTEPs to forward traffic to each other.

Spine switches in an EVPN VXLAN network are part of the underlay network and transport the VXLAN-encapsulated packets. When deployed as border nodes, spine switches connect the network with an external network and allow movement of traffic. In a BGP EVPN VXLAN fabric, spine switches can also be deployed as route reflectors.

Deployment Scenarios for Spine Switches and Leaf Switches in a BGP EVPN VXLAN Fabric

Spine switches and leaf switches in a BGP EVPN VXLAN fabric can be deployed in the following ways:

- Spine Switches and Leaf Switches in the Same Autonomous System
- Spine Switches in One Autonomous System and the Leaf Switches in a Different Autonomous System
- Spine Switches in One Autonomous System and Each Leaf Switch in a Different Autonomous System

Spine Switches and Leaf Switches in the Same Autonomous System

In this scenario, all the devices in the EVPN VXLAN network are in the same autonomous system. The spine switches function as BGP route reflectors and anycast rendezvous points (RPs). Internal Border Gateway Protocol (iBGP) is used to establish peering between the spine switches, and between the spine and leaf switches.

See [Configuration Example-Spine and Leaf Switches in the Same Autonomous System, on page 142](#) for a sample topology and configuration.

Spine Switches in One Autonomous System and the Leaf Switches in a Different Autonomous System

In this scenario, all the leaf switches are in a single autonomous system that is different from the autonomous system of the spine switches. The spine switches function as BGP route servers. iBGP is used to establish peering between the spine switches. eBGP is used to establish peering between the spine and leaf switches.

See [Configuration Example-Spine and Leaf Switches are in Different Autonomous Systems, on page 159](#) for a sample topology and configuration.

Spine Switches in One Autonomous System and Each Leaf Switch in a Different Autonomous System

In this scenario, each leaf switch is in its own individual autonomous system that is different from the autonomous system of the spine switches. The spine switches function as BGP route servers. iBGP is used to establish peering between the spine switches. eBGP is used to establish peering between the spine and leaf switches.

See [Configuration Example-Each Leaf Switch in a Different Autonomous System, on page 177](#) for a sample topology and configuration.


Configuration Examples for Spine Switches in a BGP EVPN VXLAN Network

This section provides configuration examples for spine switches for the different deployments of spine and leaf switches in a BGP EVPN VXLAN fabric.

Configuration Example-Spine and Leaf Switches in the Same Autonomous System

This section provides an example to show how spine switches are configured in a BGP EVPN VXLAN fabric using iBGP when the spine switches and leaf switches are in the same autonomous system. The example shows how to configure spine switches and verify the configuration for the topology shown below:

Figure 15: BGP EVPN VXLAN Fabric with the Spine Switches and Leaf Switches in the Same Autonomous System



The topology shows an EVPN VXLAN network with two leaf switches (VTEP 1 and VTEP 2) and two spine switches (Spine Switch 1 and Spine Switch 2). The entire BGP EVPN VXLAN fabric (which includes Spine Switch 1, Spine Switch 2, Leaf Switch 1, and Leaf Switch 2) is in autonomous system AS65001. Anycast RP is configured on both the spine switches. Spine Switch 1 and Spine Switch 2 are not route reflector clients to each other. Multicast Source Discovery Protocol (MSDP) is configured between Spine Switch 1 and Spine Switch 2 for source synchronization. Protocol Independent Multicast (PIM) is enabled on the interfaces that connect leaf switches and spine switches. Static RP is configured in the network and the underlay network uses multicast forwarding mechanism to forward BUM traffic.

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

Table 14: Configuring Spine Switch 1 and Spine Switch 2 using iBGP when the Spine Switches and the Leaf Switches are in the same Autonomous System

Spine Switch 1	Spine Switch 2
<pre> Spine-01# show running-config hostname Spine-01 ! ip routing ! ip multicast-routing ! 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 ! router ospf 1 router-id 172.16.255.1 ! router bgp 65001 template peer-policy RR-PP route-reflector-client send-community both exit-peer-policy ! template peer-session RR-PS remote-as 65001 update-source Loopback0 exit-peer-session ! bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.3 inherit peer-session RR-PS neighbor 172.16.255.4 inherit peer-session RR-PS ! address-family ipv4 exit-address-family ! </pre>	<pre> Spine-02# show running-config hostname Spine-02 ! ip routing ! ip multicast-routing ! 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 ! router ospf 1 router-id 172.16.255.2 ! router bgp 65001 template peer-policy RR-PP route-reflector-client send-community both exit-peer-policy ! template peer-session RR-PS remote-as 65001 update-source Loopback0 exit-peer-session ! bgp router-id 172.16.255.2 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.3 inherit peer-session RR-PS neighbor 172.16.255.4 inherit peer-session RR-PS ! address-family ipv4 exit-address-family ! </pre>

Spine Switch 1	Spine Switch 2
<pre>address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community extended neighbor 172.16.255.3 inherit peer-policy RR-PP neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community extended neighbor 172.16.255.4 inherit peer-policy RR-PP 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#</pre>	<pre>address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community extended neighbor 172.16.255.3 inherit peer-policy RR-PP neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community extended neighbor 172.16.255.4 inherit peer-policy RR-PP 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#</pre>

Table 15: Configuring Leaf Switch 1 and Leaf Switch 2 using iBGP when the Spine Switches and the Leaf Switches are in the same Autonomous System

Leaf Switch 1	Leaf Switch 2
<pre> 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 ! 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 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 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 ! </pre>	<pre> 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 ! 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 replication-type ingress ! 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 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 ! </pre>

Leaf Switch 1	Leaf Switch 2
<pre> 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 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 mcast-group 225.0.0.101 member vni 10102 ingress-replication member vni 50901 vrf green ! 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 ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family ! </pre>	<pre> 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 mode trunk ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 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 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 ! </pre>

Leaf Switch 1	Leaf Switch 2
<pre> address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-01# </pre>	<pre> address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-02# </pre>

The following examples provide sample outputs for **show** commands on the devices in the topology configured in the preceding tables:

Spine Switch 1

The following example shows the output for the **show ip ospf neighbor** command on spine switch 1:

```
Spine-01# show ip ospf neighbor
```

Neighbor ID	Pri	State		Dead Time	Address	Interface
172.16.255.4	0	FULL/	-	00:00:39	172.16.14.4	GigabitEthernet1/0/2
172.16.255.3	0	FULL/	-	00:00:30	172.16.13.3	GigabitEthernet1/0/1

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 46, main routing table version 46
18 network entries using 6192 bytes of memory
38 path entries using 7904 bytes of memory
14/13 BGP path/bestpath attribute entries using 4032 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
12 BGP extended community entries using 640 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 18848 total bytes of memory
BGP activity 27/9 prefixes, 49/11 paths, scan interval 60 secs
18 networks peaked at 17:16:59 May 24 2020 UTC (22:49:24.588 ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.3	4	65001	1517	1536	46	0	0	22:49:32	9
172.16.255.4	4	65001	1297	1310	46	0	0	19:23:05	11

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 1 for route type 2 and the IP address of host device 1:

```
Spine-01# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
```

```

BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 4

```

```
Paths: (2 available, best #2, table EVPN-BGP-Table)
```

```
Advertised to update-groups:
```

```
1 2
```

```

Refresh Epoch 1
Local
 172.16.254.3 (metric 2) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
  Router MAC:10B3.D56A.8FC8
  Originator: 172.16.255.3, Cluster list: 172.16.255.2
  rx pathid: 0, tx pathid: 0
  net: 0x7F54CCA547D0, path: 0x7F54CCA63D70, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on May 24 2020 20:42:55 UTC
Refresh Epoch 2
Local, (Received from a RR-client)
 172.16.254.3 (metric 2) (via default) from 172.16.255.3 (172.16.255.3)
  Origin incomplete, metric 0, localpref 100, valid, internal, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
  Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F54CCA547D0, path: 0x7F54CCA64AF0, pathext: 0x7F54CA789BA8
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 24 2020 17:16:50 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 1 for route type 2 and the IP address of host device 2:

```

Spine-01# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
 version 42
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1      2
Refresh Epoch 2
Local, (Received from a RR-client)
 172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
  Origin incomplete, metric 0, localpref 100, valid, internal, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
  Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F54CCA53E30, path: 0x7F54CCA63428, pathext: 0x7F54CA7898A8
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 24 2020 20:43:18 UTC
Refresh Epoch 1
Local
 172.16.254.4 (metric 2) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
  Router MAC:7C21.0DBD.9548
  Originator: 172.16.255.4, Cluster list: 172.16.255.2
  rx pathid: 0, tx pathid: 0
  net: 0x7F54CCA53E30, path: 0x7F54CCA64280, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on May 24 2020 20:28:04 UTC

```

The following example shows the output for the **show ip pim neighbor** command on spine switch 1:

```

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                               1d22h/00:01:41      v2       1 / DR S P G
172.16.13.3   GigabitEthernet1/0/1
172.16.14.4   GigabitEthernet1/0/2   4w5d/00:01:24      v2       1 / DR S P G

```

The following example shows the output for the **show ip pim rp map** command on spine switch 1:

```

Spine-01# show ip pim rp map
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
          RP: 172.16.255.255 (?)

```

The following example shows the output for the **show ip rpf** command on spine switch 1:

```

Spine-01# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: Loopback2
  RPF neighbor: ? (172.16.255.255) - directly connected
  RPF route/mask: 172.16.255.255/32
  RPF type: multicast (connected)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base

```

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

```

Spine-01# show ip msdp summary
MSDP Peer Status Summary
Peer Address      AS      State      Uptime/   Reset SA   Peer Name
                  AS              Downtime  Count     Count
172.16.254.2     65001  Up         22:37:35  0         2         ?

```

The following example shows the output for the **show ip msdp sa-cache** command on spine switch 1:

```

Spine-01# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(172.16.254.3, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:00:29/00:05:30, Peer 172.16.254.2
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:00:17/00:05:43, Peer 172.16.254.2

```

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

```

Spine-01# show ip mroute 225.0.0.10
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

(*, 225.0.0.101), 00:01:04/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), 00:00:51/00:02:08, 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), 00:01:04/00:01:55, flags: PA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.13.3
  Outgoing interface list: Null

```

Spine Switch 2

The following example shows the output for the **show ip ospf neighbor** command on spine switch 2:

```
Spine-02# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.4	0	FULL/ -	00:00:39	172.16.24.4	GigabitEthernet1/0/2
172.16.255.3	0	FULL/ -	00:00:35	172.16.23.3	GigabitEthernet1/0/1

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 28, main routing table version 28
18 network entries using 6192 bytes of memory
38 path entries using 7904 bytes of memory
14/13 BGP path/bestpath attribute entries using 4032 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
12 BGP extended community entries using 640 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 18848 total bytes of memory
BGP activity 36/18 prefixes, 58/20 paths, scan interval 60 secs
18 networks peaked at 16:03:20 May 24 2020 UTC (1d00h ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.3	4	65001	1307	1322	28	0	0	19:35:35	9
172.16.255.4	4	65001	1316	1334	28	0	0	19:51:36	11

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 2 for route type 2 and the IP address of host device 1:

```
Spine-02# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
```

```

BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 24
Paths: (2 available, best #1, table EVPN-BGP-Table)

```

```

Advertised to update-groups:
  2          3
Refresh Epoch 2
Local, (Received from a RR-client)
 172.16.254.3 (metric 2) (via default) from 172.16.255.3 (172.16.255.3)
  Origin incomplete, metric 0, localpref 100, valid, internal, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0x0
  net: 0x7FEFE69D6638, path: 0x7FEFE45FED18, pathext: 0x7FEFE6645CC0
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 24 2020 20:43:24 UTC
Refresh Epoch 1
Local
 172.16.254.3 (metric 2) (via default) from 172.16.255.1 (172.16.255.1)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  Originator: 172.16.255.3, Cluster list: 172.16.255.1
  rx pathid: 0, tx pathid: 0
  net: 0x7FEFE69D6638, path: 0x7FEFE45FF738, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on May 24 2020 20:27:33 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 2 for route type 2 and the IP address of host device 2:

```

Spine-02# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 10
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    2          3
  Refresh Epoch 1
  Local
    172.16.254.4 (metric 2) (via default) from 172.16.255.1 (172.16.255.1)
    Origin incomplete, metric 0, localpref 100, valid, internal
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
    Originator: 172.16.255.4, Cluster list: 172.16.255.1
    rx pathid: 0, tx pathid: 0
    net: 0x7FEFE69D64D8, path: 0x7FEFE45FE730, pathext: 0x0
    flags: net: 0x0, path: 0x3, pathext: 0x0
    Updated on May 24 2020 20:43:46 UTC
  Refresh Epoch 1
  Local, (Received from a RR-client)
    172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
    rx pathid: 0, tx pathid: 0x0
    net: 0x7FEFE69D64D8, path: 0x7FEFE45FF660, pathext: 0x7FEFE6645B40
    flags: net: 0x0, path: 0x3, pathext: 0x81
    Updated on May 24 2020 20:27:22 UTC

```

The following example shows the output for the **show ip pim neighbor** command on spine switch 2:

```
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
172.16.23.3       GigabitEthernet1/0/1     6w3d/00:01:21    v2    1 / DR S P G
172.16.24.4       GigabitEthernet1/0/2     1d22h/00:01:18    v2    1 / DR S P G
```

The following example shows the output for the **show ip pim rp map** command on spine switch 2:

```
Spine-02# show ip pim rp map
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
          RP: 172.16.255.255 (?)
```

The following example shows the output for the **show ip rpf** command on spine switch 2:

```
Spine-02# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: Loopback2
  RPF neighbor: ? (172.16.255.255) - directly connected
  RPF route/mask: 172.16.255.255/32
  RPF type: multicast (connected)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base
```

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

```
Spine-02# 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       22:41:13 3        2        ?
```

The following example shows the output for the **show ip msdp sa-cache** command on spine switch 2:

```
Spine-02# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(172.16.254.3, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:04:09/00:05:57, Peer 172.16.254.1
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:03:56/00:05:57, Peer 172.16.254.1
```

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

```
Spine-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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 5w6d/00:03:16, RP 172.16.255.255, flags: S
Incoming interface: Null, RPF nbr 0.0.0.0
Outgoing interface list:
  GigabitEthernet1/0/2, Forward/Sparse, 1d22h/00:03:10
  GigabitEthernet1/0/1, Forward/Sparse, 5w6d/00:02:55

(172.16.254.4, 225.0.0.101), 00:00:13/00:02:46, flags: TA
Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
Outgoing interface list:
  GigabitEthernet1/0/1, Forward/Sparse, 00:00:13/00:03:16

(172.16.254.3, 225.0.0.101), 00:00:23/00:02:36, flags: A
Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
Outgoing interface list:
  GigabitEthernet1/0/2, Forward/Sparse, 00:00:23/00:03:10

```

Leaf Switch 1

The following example shows the output for the **show ip ospf neighbor** command on leaf switch 1:

```
Leaf-01# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.2	0	FULL/ -	00:00:34	172.16.23.2	GigabitEthernet1/0/2
172.16.255.1	0	FULL/ -	00:00:30	172.16.13.1	GigabitEthernet1/0/1

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

```
Leaf-01# show bgp l2vpn evpn summary
```

```

BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 11429, main routing table version 11429
27 network entries using 9288 bytes of memory
36 path entries using 7488 bytes of memory
15/15 BGP path/bestpath attribute entries using 4320 bytes of memory
2 BGP rinfo entries using 80 bytes of memory
12 BGP extended community entries using 624 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 21800 total bytes of memory
BGP activity 398/365 prefixes, 4243/4201 paths, scan interval 60 secs
89 networks peaked at 20:32:14 Apr 21 2020 UTC (4w5d ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	261	242	11429	0	0	03:28:13	9
172.16.255.2	4	65001	31	16	11429	0	0	00:02:08	9

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 host device 2:


```

Leaf-01# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 11423
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  Local, imported path from [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24
(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
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F575DB9FAB0, path: 0x7F575FD77698, pathext: 0x7F575DBD5B48, exp_net:
0x7F575DBA3B50
      flags: net: 0x0, path: 0x40000000000003, pathext: 0x81
      Updated on May 24 2020 20:40:59 UTC
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 11414
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
      EVPN ESI: 000000000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      net: 0x7F575DBA3B50, path: 0x7F575FD77E30, pathext: 0x0
      flags: net: 0x0, path: 0x3, pathext: 0x0
      Updated on May 24 2020 20:40:37 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
      EVPN ESI: 000000000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F575DBA3B50, path: 0x7F575FD769F0, pathext: 0x7F575DBD5D88
      flags: net: 0x0, path: 0x3, pathext: 0x81
      Updated on May 24 2020 20:40:59 UTC

```

The following example shows the output for the **show ip pim neighbor** command on leaf switch 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
172.16.13.1       GigabitEthernet1/0/1     1d03h/00:01:21    v2   1 / S P G
172.16.23.2       GigabitEthernet1/0/2     6w2d/00:01:25    v2   1 / S P G

```

The following example shows the output for the **show ip pim rp mapping** command on leaf switch 1:

```
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 (?)
```

The following example shows the output for the **show ip ro** command on leaf switch 1:

```
Leaf-01# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 172.16.13.1 on GigabitEthernet1/0/1, 1d03h ago
Routing Descriptor Blocks:
  * 172.16.23.2, from 172.16.255.2, 4w5d ago, via GigabitEthernet1/0/2
    Route metric is 2, traffic share count is 1
  172.16.13.1, from 172.16.255.1, 1d03h ago, via GigabitEthernet1/0/1
    Route metric is 2, traffic share count is 1
```

The following example shows the output for the **show ip rpf** command on leaf switch 1:

```
Leaf-01# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: GigabitEthernet1/0/2
  RPF neighbor: ? (172.16.23.2)
  RPF route/mask: 172.16.255.255/32
  RPF type: unicast (ospf 1)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

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

```
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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 6w2d/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, 6w2d/00:01:57

(172.16.254.4, 225.0.0.101), 00:00:49/00:02:10, flags: JTx
```

```

Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
Outgoing interface list:
  Tunnel0, Forward/Sparse-Dense, 00:00:49/00:02:10
(172.16.254.3, 225.0.0.101), 00:01:01/00:01:58, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:01/00:03:27

```

Leaf Switch 2

The following example shows the output for the **show ip ospf neighbor** command on leaf switch 2:

```
Leaf-02# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.2	0	FULL/ -	00:00:34	172.16.24.2	GigabitEthernet1/0/2
172.16.255.1	0	FULL/ -	00:00:35	172.16.14.1	GigabitEthernet1/0/1

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

```
Leaf-02# show bgp l2vpn evpn summary
```

```

BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 168, main routing table version 168
25 network entries using 8600 bytes of memory
36 path entries using 7488 bytes of memory
16/15 BGP path/bestpath attribute entries using 4608 bytes of memory
2 BGP rrinfo entries using 80 bytes of memory
13 BGP extended community entries using 664 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 21440 total bytes of memory
BGP activity 70/39 prefixes, 168/124 paths, scan interval 60 secs
31 networks peaked at 15:56:08 May 24 2020 UTC (05:05:36.264 ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	45	31	168	0	0	00:16:18	9
172.16.255.2	4	65001	54	48	168	0	0	00:32:42	9

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 host device 1:

```
Leaf-02# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
```

```

BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 163
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 2
  Local
    172.16.254.3 (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, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
      Originator: 172.16.255.3, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F84B9145020, path: 0x7F84BB3355F8, pathext: 0x7F84BB5B4318

```

```

    flags: net: 0x0, path: 0x3, pathext: 0x81
    Updated on May 24 2020 20:45:25 UTC
Refresh Epoch 1
Local
  172.16.254.3 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
    Origin incomplete, metric 0, localpref 100, valid, internal
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
    Originator: 172.16.255.3, Cluster list: 172.16.255.2
    rx pathid: 0, tx pathid: 0
    net: 0x7F84B9145020, path: 0x7F84BB333948, pathext: 0x0
    flags: net: 0x0, path: 0x3, pathext: 0x0
    Updated on May 24 2020 20:45:03 UTC
BGP routing table entry for [2][172.16.254.4:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 166
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 2
  Local, imported path from [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
(global)
    172.16.254.3 (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, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
        Router MAC:10B3.D56A.8FC8
      Originator: 172.16.255.3, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F84B9145700, path: 0x7F84BB334008, pathext: 0x7F84BB5B3A18, exp_net:
0x7F84B9145020
      flags: net: 0x0, path: 0x40000000000003, pathext: 0x81
      Updated on May 24 2020 20:45:25 UTC

```

The following example shows the output for the **show ip pim neighbor** command on leaf switch 2:

```

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
172.16.14.1       GigabitEthernet1/0/1     4w5d/00:01:26    v2   1 / S P G
172.16.24.2       GigabitEthernet1/0/2     1d03h/00:01:20    v2   1 / S P G

```

The following example shows the output for the **show ip pim rp map** command on leaf switch 2:

```

Leaf-02# show ip pim rp map
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
          RP: 172.16.255.255 (?)

```

The following example shows the output for the **show ip ro** command on leaf switch 2:

```

Leaf-02# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 172.16.14.1 on GigabitEthernet1/0/1, 05:12:11 ago

```

```

Routing Descriptor Blocks:
* 172.16.24.2, from 172.16.255.2, 05:12:11 ago, via GigabitEthernet1/0/2
  Route metric is 2, traffic share count is 1
  172.16.14.1, from 172.16.255.1, 05:12:11 ago, via GigabitEthernet1/0/1
    Route metric is 2, traffic share count is 1

```

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

```

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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 3d07h/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, 3d07h/00:00:38

(172.16.254.4, 225.0.0.101), 00:00:09/00:02:50, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:00:09/00:03:20


(172.16.254.3, 225.0.0.101), 00:00:28/00:02:31, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:28/00:02:31

```

Configuration Example-Spine and Leaf Switches are in Different Autonomous Systems

This section provides an example to show how spine switches are configured in a BGP EVPN VXLAN fabric using eBGP when the spine switches are in one autonomous system and the leaf switches are in a different autonomous system. The example shows how to configure spine switches and verify the configuration for the topology shown below:

Figure 16: BGP EVPN VXLAN Fabric with the Spine Switches in One Autonomous System and the Leaf Switches in a Different Autonomous System



The topology shows an EVPN VXLAN network with two leaf switches (Leaf Switch 1 and Leaf Switch 2) and two spine switches (Spine Switch 1 and Spine Switch 2). Spine Switch 1 and Spine Switch 2 are in autonomous system AS65001. Leaf Switch 1 and Leaf Switch 2 are in autonomous system AS65002. Spine Switch 1 and Spine Switch 2 are BGP route servers and are not route reflector clients to each other. Multicast Source Discovery Protocol (MSDP) is configured between Spine Switch 1 and Spine Switch 2 for source synchronization. Protocol Independent Multicast (PIM) is enabled on the interfaces that connect leaf switches and spine switches. Static RP is configured in the network and the underlay network uses multicast forwarding mechanism to forward BUM traffic.



Note

- You must run the **neighbor ip-address allow-as-in** command in the L2VPN EVPN address family configuration mode on the leaf switches to allow processing of BGP updates that have a different autonomous system number.
- You must manually run the **no bgp default route-target filter** command in router configuration mode on the spine switches.
- You must configure eBGP multihop on the leaf and spine switches for the fabric to function.

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

Table 16: Configuring Spine Switch 1 and Spine Switch 2 using eBGP when the Spine Switches are in one Autonomous System and the Leaf Switches are in a Different Autonomous System

Spine Switch 1	Spine Switch 2
<pre> Spine-01# show running-config hostname Spine-01 ! ip routing ! ip multicast-routing ! 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 ! 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 no bgp default route-target filter neighbor 172.16.255.3 remote-as 65002 neighbor 172.16.255.3 ebgp-multihop 255 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65002 neighbor 172.16.255.4 ebgp-multihop 255 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family ! </pre>	<pre> Spine-02# show running-config hostname Spine-02 ! ip routing ! ip multicast-routing ! 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 ! 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 no bgp default route-target filter neighbor 172.16.255.3 remote-as 65002 neighbor 172.16.255.3 ebgp-multihop 255 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65002 neighbor 172.16.255.4 ebgp-multihop 255 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family ! </pre>

Spine Switch 1	Spine Switch 2
<pre> address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community extended neighbor 172.16.255.3 route-map BGP-NHU out neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-map BGP-NHU out 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 ! route-map BGP-NHU permit 10 set ip next-hop unchanged ! end Spine-01# </pre>	<pre> address-family l2vpn evpn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-map BGP-NHU out neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-map BGP-NHU out 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 ! route-map BGP-NHU permit 10 set ip next-hop unchanged ! end Spine-02# </pre>

Table 17: Configuring Leaf Switch 1 and Leaf Switch 2 using eBGP when the Spine Switches are in one Autonomous System and the Leaf Switches are in a Different Autonomous System

Leaf Switch 1	Leaf Switch 2
<pre> 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 ! 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 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 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 </pre>	<pre> 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 ! 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 replication-type ingress ! 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 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 ! </pre>

Leaf Switch 1	Leaf Switch 2
<pre> 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 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 mcast-group 225.0.0.101 member vni 10102 ingress-replication member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.3 ! router bgp 65002 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 ebgp-multihop 255 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 ebgp-multihop 255 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 allowas-in neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.2 allowas-in exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family </pre>	<pre> 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 mode trunk ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.4 ! router bgp 65002 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 ebgp-multihop 255 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 ebgp-multihop 255 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 allowas-in neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both neighbor 172.16.255.2 allowas-in exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected redistribute static exit-address-family </pre>

Leaf Switch 1	Leaf Switch 2
<pre>! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-01#</pre>	<pre>! address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-02#</pre>

The following examples provide sample outputs for **show** commands on the devices in the topology configured in the preceding tables:

Spine Switch 1

The following example shows the output for the **show ip ospf neighbor** command on spine switch 1:

```
Spine-01# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.4	0	FULL/ -	00:00:33	172.16.14.4	GigabitEthernet1/0/2
172.16.255.3	0	FULL/ -	00:00:34	172.16.13.3	GigabitEthernet1/0/1

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 75, main routing table version 75
18 network entries using 6192 bytes of memory
38 path entries using 7904 bytes of memory
27/13 BGP path/bestpath attribute entries using 7776 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
12 BGP extended community entries using 640 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 22536 total bytes of memory
BGP activity 18/0 prefixes, 76/38 paths, scan interval 60 secs
18 networks peaked at 20:34:25 May 27 2020 UTC (5d18h ago)

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.255.3  4      65002   8446   8456     75   0    0 5d07h     9
172.16.255.4  4      65002   8446   8447     75   0    0 5d07h    11
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 1 for route type 2 and the IP address of host device 2:

```
Spine-01# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 72
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    4          5
```

```

Refresh Epoch 2
65002
 172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
  Origin incomplete, metric 0, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
    Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F789AD67240, path: 0x7F789AD76820, pathext: 0x7F789AD88298
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 28 2020 07:29:30 UTC
Refresh Epoch 1
65002
 172.16.254.4 (metric 2) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
    Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0
  net: 0x7F789AD67240, path: 0x7F789AD76EE0, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on May 28 2020 07:27:54 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 1 for route type 2 and the IP address of host device 2:

```

Spine-01# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 40
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    4          5
Refresh Epoch 1
65002
 172.16.254.3 (metric 2) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0
  net: 0x7F789AD67EA0, path: 0x7F789AD77678, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on May 28 2020 07:29:03 UTC
Refresh Epoch 1
65002
 172.16.254.3 (metric 2) (via default) from 172.16.255.3 (172.16.255.3)
  Origin incomplete, metric 0, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F789AD67EA0, path: 0x7F789AD77FC0, pathext: 0x7F789AD88598
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 28 2020 07:27:47 UTC

```

The following example shows the output for the **show ip pim neighbor** command on spine switch 1:

```

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
172.16.13.3   GigabitEthernet1/0/1   5d19h/00:01:44   v2    1 / DR S P G
172.16.14.4   GigabitEthernet1/0/2   5d19h/00:01:36   v2    1 / DR S P G

```

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 (?)

```

The following example shows the output for the **show ip ro** command on spine switch 1:

```

Spine-01# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Loopback2
    Route metric is 0, traffic share count is 1

```

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

```

Spine-01# 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         5d19h    0      2      ?

```

The following example shows the output for the **show ip msdp sa-cache** command on spine switch 1:

```

Spine-01# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(172.16.254.3, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:04:01/00:05:23, Peer 172.16.254.2
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:03:39/00:05:26, Peer 172.16.254.2

```

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

```

Spine-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-Route 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

(*, 225.0.0.101), 00:04:02/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), 00:00:34/00:02:25, 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), 00:00:46/00:02:13, flags: PA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.13.3
  Outgoing interface list: Null

```

Spine Switch 2

The following example shows the output for the **show ip ospf neighbor** command on spine switch 2:

```
Spine-02# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.4	0	FULL/ -	00:00:37	172.16.24.4	GigabitEthernet1/0/2
172.16.255.3	0	FULL/ -	00:00:32	172.16.23.3	GigabitEthernet1/0/1

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 91, main routing table version 91
18 network entries using 6192 bytes of memory
38 path entries using 7904 bytes of memory
27/13 BGP path/bestpath attribute entries using 7776 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
12 BGP extended community entries using 640 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 22536 total bytes of memory
BGP activity 20/2 prefixes, 76/38 paths, scan interval 60 secs
18 networks peaked at 20:36:02 May 27 2020 UTC (5d18h ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.3	4	65002	8443	8442	91	0	0	5d07h	9
172.16.255.4	4	65002	8442	8446	91	0	0	5d07h	11

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 2 for host device 1:

```
Spine-02# bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
```

```

BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 74
Paths: (2 available, best #2, table EVPN-BGP-Table)

```

```

Advertised to update-groups:
  3          4
Refresh Epoch 1
65002
  172.16.254.4 (metric 2) (via default) from 172.16.255.1 (172.16.255.1)
    Origin incomplete, metric 0, localpref 100, valid, internal
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
    rx pathid: 0, tx pathid: 0
    net: 0x7FB64B5D07C0, path: 0x7FB64B5DFA08, pathext: 0x0
    flags: net: 0x0, path: 0x3, pathext: 0x0
    Updated on May 28 2020 07:30:01 UTC
Refresh Epoch 1
65002
  172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
    Origin incomplete, metric 0, localpref 100, valid, external, best
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
    rx pathid: 0, tx pathid: 0x0
    net: 0x7FB64B5D07C0, path: 0x7FB64B5E01A0, pathext: 0x7FB64B5F1498
    flags: net: 0x0, path: 0x3, pathext: 0x81
    Updated on May 28 2020 07:28:25 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 2 for host device 2:

```

Spine-02# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 88
Paths: (2 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    3          4
Refresh Epoch 2
65002
  172.16.254.3 (metric 2) (via default) from 172.16.255.3 (172.16.255.3)
    Origin incomplete, metric 0, localpref 100, valid, external, best
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
    rx pathid: 0, tx pathid: 0x0
    net: 0x7FB64B5D1580, path: 0x7FB64B5E0D70, pathext: 0x7FB64B5F19D8
    flags: net: 0x0, path: 0x3, pathext: 0x81
    Updated on May 28 2020 07:29:33 UTC
Refresh Epoch 1
65002
  172.16.254.3 (metric 2) (via default) from 172.16.255.1 (172.16.255.1)
    Origin incomplete, metric 0, localpref 100, valid, internal
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
    rx pathid: 0, tx pathid: 0
    net: 0x7FB64B5D1580, path: 0x7FB64B5E0AE8, pathext: 0x0
    flags: net: 0x0, path: 0x3, pathext: 0x0
    Updated on May 28 2020 07:28:18 UTC

```

The following example shows the output for the **show ip pim neighbor** command on spine switch 2:

```

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                               5d19h/00:01:33   v2    Prio/Mode
172.16.23.3   GigabitEthernet1/0/1
172.16.24.4   GigabitEthernet1/0/2   5d19h/00:01:18   v2    1 / DR S P G

```

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 (?)

```

The following example shows the output for the **show ip ro** command on spine switch 2:

```

Spine-02# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Loopback2
    Route metric is 0, traffic share count is 1

```

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

```

Spine-02# 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         5d19h    0      2      ?

```

The following example shows the output for the **show ip msdp sa-cache** command on spine switch 2:

```

Spine-02# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(172.16.254.3, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:04:07/00:05:17, Peer 172.16.254.1
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:03:45/00:05:20, Peer 172.16.254.1

```

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

```

Spine-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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 5d19h/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, 5d19h/00:03:15
    GigabitEthernet1/0/1, Forward/Sparse, 5d19h/00:03:21

(172.16.254.4, 225.0.0.101), 00:00:40/00:02:19, flags: A
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:00:40/00:03:21

(172.16.254.3, 225.0.0.101), 00:00:52/00:02:07, flags: A
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:00:52/00:03:15

```

Leaf Switch 1

The following example shows the output for the **show ip ospf neighbor** command on leaf switch 1:

```
Leaf-01# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.2	0	FULL/ -	00:00:38	172.16.23.2	GigabitEthernet1/0/2
172.16.255.1	0	FULL/ -	00:00:32	172.16.13.1	GigabitEthernet1/0/1

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

```
Leaf-01# show bgp l2vpn evpn summary
```

```

BGP router identifier 172.16.255.3, local AS number 65002
BGP table version is 32, main routing table version 32
27 network entries using 9288 bytes of memory
38 path entries using 7904 bytes of memory
16/15 BGP path/bestpath attribute entries using 4608 bytes of memory
1 BGP AS-PATH entries using 40 bytes of memory
13 BGP extended community entries using 664 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 22504 total bytes of memory
BGP activity 395/362 prefixes, 918/872 paths, scan interval 60 secs
27 networks peaked at 13:15:47 May 26 2020 UTC (1w0d ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	8457	8446	32	0	0	5d07h	9
172.16.255.2	4	65001	8443	8444	32	0	0	5d07h	11

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 host device 2:

```
Leaf-01# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
```

```

BGP routing table entry for [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 22

```

```

Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  65001 65002, imported path from
[2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24 (global)
  172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
  Origin incomplete, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
  Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F575E47B150, path: 0x7F575E1EF800, pathext: 0x7F575E201C08, exp_net:
0x7F575E479470
  flags: net: 0x0, path: 0x40000000000003, pathext: 0x81
  Updated on May 28 2020 07:25:32 UTC
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 10
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
  19
  Refresh Epoch 2
  65001 65002
  172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, localpref 100, valid, external
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
  Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0
  net: 0x7F575E479470, path: 0x7F575E1EFD10, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on May 28 2020 07:26:48 UTC
  Refresh Epoch 1
  65001 65002
  172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
  Origin incomplete, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
  Router MAC:7C21.0DBD.9548

```

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 host device 1:

```

Leaf-01# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 4
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
  19
  Refresh Epoch 1
  Local
  :: (via default) from 0.0.0.0 (172.16.255.3)
  Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65002:101 ENCAP:8
  Router MAC:10B3.D56A.8FC8
  Local irb vxlan vtep:
  vrf:green, l3-vni:50901
  local router mac:10B3.D56A.8FC8
  core-irb interface:Vlan901
  vtep-ip:172.16.254.3
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F575E479B50, path: 0x7F575E1F0580, pathext: 0x7F575E201CC8

```

```
flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
Updated on May 28 2020 07:25:30 UTC
```

The following example shows the output for the **show ip pim neighbor** command on leaf switch 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
172.16.13.1       GigabitEthernet1/0/1     5d19h/00:01:38   v2    1 / S P G
172.16.23.2       GigabitEthernet1/0/2     5d19h/00:01:17   v2    1 / S P G
```

The following example shows the output for the **show ip pim rp mapping** command on leaf switch 1:

```
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 (?)
```

The following example shows the output for the **show ip ro** command on leaf switch 1:

```
Leaf-01# show ip ro 172.16.255.255 Routing entry for 172.16.255.255/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 172.16.23.2 on GigabitEthernet1/0/2, 5d19h ago
  Routing Descriptor Blocks:
    172.16.23.2, from 172.16.255.2, 5d19h ago, via GigabitEthernet1/0/2
      Route metric is 2, traffic share count is 1
    * 172.16.13.1, from 172.16.255.1, 5d19h ago, via GigabitEthernet1/0/1
      Route metric is 2, traffic share count is 1
```

The following example shows the output for the **show ip rpf** command on leaf switch 1:

```
Leaf-01# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: GigabitEthernet1/0/2
  RPF neighbor: ? (172.16.23.2)
  RPF route/mask: 172.16.255.255/32
  RPF type: unicast (ospf 1)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

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

```
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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 7w4d/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, 1w0d/00:00:40

(172.16.254.4, 225.0.0.101), 00:01:22/00:01:37, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:01:22/00:01:37

(172.16.254.3, 225.0.0.101), 00:01:35/00:01:24, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:35/00:02:53

```

Leaf Switch 2

The following example shows the output for the **show ip ospf neighbor** command on leaf switch 2:

```
Leaf-02# show ip ospf neighbor
```

Neighbor ID	Pri	State	Dead Time	Address	Interface
172.16.255.2	0	FULL/ -	00:00:34	172.16.24.2	GigabitEthernet1/0/2
172.16.255.1	0	FULL/ -	00:00:31	172.16.14.1	GigabitEthernet1/0/1

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

```
Leaf-02# show bgp l2vpn evpn summary
```

```

BGP router identifier 172.16.255.4, local AS number 65002
BGP table version is 28, main routing table version 28
25 network entries using 8600 bytes of memory
34 path entries using 7072 bytes of memory
16/15 BGP path/bestpath attribute entries using 4608 bytes of memory
1 BGP AS-PATH entries using 40 bytes of memory
13 BGP extended community entries using 664 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 20984 total bytes of memory
BGP activity 199/168 prefixes, 638/596 paths, scan interval 60 secs
25 networks peaked at 13:20:44 May 26 2020 UTC (1w0d ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	8449	8447	28	0	0	5d07h	9
172.16.255.2	4	65001	8448	8443	28	0	0	5d07h	7

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 host device 1:

```
Leaf-02# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
```

```

BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 4

```

```

Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    7
  Refresh Epoch 2
  65001 65002
    172.16.254.3 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, localpref 100, valid, external
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
      rx pathid: 0, tx pathid: 0
      net: 0x7F84BB3C4290, path: 0x7F84BB49BF98, pathext: 0x0
      flags: net: 0x0, path: 0x3, pathext: 0x0
      Updated on May 28 2020 07:31:42 UTC
  Refresh Epoch 1
  65001 65002
    172.16.254.3 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, localpref 100, valid, external, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F84BB3C4290, path: 0x7F84BB49D9C0, pathext: 0x7F84BB594138
      flags: net: 0x0, path: 0x3, pathext: 0x81
      Updated on May 28 2020 07:31:37 UTC
BGP routing table entry for [2][172.16.254.4:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
  version 20
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  65001 65002, imported path from
  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24 (global)
    172.16.254.3 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, localpref 100, valid, external, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8

```

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 host device 2:

```

Leaf-02# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
  version 10
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
    7
  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
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
    Router MAC:7C21.0DBD.9548
  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
    net: 0x7F84BB3C4970, path: 0x7F84BB49CDF0, pathext: 0x7F84BB593CB8
    flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
    Updated on May 28 2020 07:30:04 UTC

```

The following example shows the output for the **show ip pim neighbor** command on leaf switch 2:

```
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
172.16.14.1       GigabitEthernet1/0/1  5d19h/00:01:22   v2    1 / S P G
172.16.24.2       GigabitEthernet1/0/2  5d19h/00:01:27   v2    1 / S P G
```

The following example shows the output for the **show ip pim rp mapping** command on leaf switch 2:

```
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 (?)
```

The following example shows the output for the **show ip ro** command on leaf switch 2:

```
Leaf-02# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 172.16.24.2 on GigabitEthernet1/0/2, 5d19h ago
  Routing Descriptor Blocks:
    172.16.24.2, from 172.16.255.2, 5d19h ago, via GigabitEthernet1/0/2
      Route metric is 2, traffic share count is 1
    * 172.16.14.1, from 172.16.255.1, 5d19h ago, via GigabitEthernet1/0/1
      Route metric is 2, traffic share count is 1
```

The following example shows the output for the **show ip rpf** command on leaf switch 2:

```
Leaf-02# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: GigabitEthernet1/0/2
  RPF neighbor: ? (172.16.24.2)
  RPF route/mask: 172.16.255.255/32
  RPF type: unicast (ospf 1)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

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

```
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
```

```
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 1w5d/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, 1w5d/00:00:06


(172.16.254.4, 225.0.0.101), 00:01:56/00:01:03, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:56/00:02:32

(172.16.254.3, 225.0.0.101), 00:02:09/00:00:50, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:02:09/00:00:50
```

Configuration Example-Each Leaf Switch in a Different Autonomous System

This section provides an example to show how spine switches are configured in a BGP EVPN VXLAN fabric using eBGP when the spine switches are in one autonomous system and each leaf switch is in a different autonomous system. The example shows how to configure spine switches and verify the configuration for the topology shown below:

Figure 17: BGP EVPN VXLAN Fabric with the Spine Switches in one Autonomous System and each Leaf Switch in a Different Autonomous System



The topology shows an EVPN VXLAN network with two leaf switches (Leaf Switch 1 and Leaf Switch 2) and two spine switches (Spine Switch 1 and Spine Switch 2). Spine Switch 1 and Spine Switch 2 are in autonomous system AS65001. Leaf Switch 1 is in autonomous system AS65002. Leaf Switch 2 is in autonomous system AS65003. Spine Switch 1 and Spine Switch 2 are BGP route servers and are not route reflector clients to each other. Multicast Source Discovery Protocol (MSDP) is configured between Spine Switch 1 and Spine Switch 2 for source synchronization. Protocol Independent Multicast (PIM) is enabled on the interfaces that connect leaf switches and spine switches. Static RP is configured in the network and the underlay network uses multicast forwarding mechanism to forward BUM traffic.



Note

- You must run the **rewrite-evpn-rt-asn** command in the L2VPN EVPN address family configuration mode on the leaf switches to allow processing of BGP updates that have a different autonomous system number.
- You must manually run the **no bgp default route-target filter** command in router configuration mode on the spine switches.
- You must configure eBGP multihop on the leaf and spine switches for the fabric to function.

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

Table 18: Configuring Spine Switch 1 and Spine Switch 2 using eBGP when the Spine Switches are in one Autonomous System and each Leaf Switch is in a Different Autonomous System

Spine Switch 1	Spine Switch 2
<pre> Spine-01# show running-config hostname Spine-01 ! ip routing ! ip multicast-routing ! 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 ! 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 no bgp default route-target filter neighbor 172.16.255.3 remote-as 65002 neighbor 172.16.255.3 ebgp-multihop 255 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65003 neighbor 172.16.255.4 ebgp-multihop 255 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family ! </pre>	<pre> Spine-02# show running-config hostname Spine-02 ! ip routing ! ip multicast-routing ! 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 ! 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 no bgp default route-target filter neighbor 172.16.255.3 remote-as 65002 neighbor 172.16.255.3 ebgp-multihop 255 neighbor 172.16.255.3 update-source Loopback0 neighbor 172.16.255.4 remote-as 65003 neighbor 172.16.255.4 ebgp-multihop 255 neighbor 172.16.255.4 update-source Loopback0 ! address-family ipv4 exit-address-family ! </pre>

Spine Switch 1	Spine Switch 2
<pre> address-family l2vpn evpn rewrite-evpn-rt-asn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-map BGP-NHU out neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-map BGP-NHU out 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 ! route-map BGP-NHU permit 10 set ip next-hop unchanged ! end Spine-01# </pre>	<pre> address-family l2vpn evpn rewrite-evpn-rt-asn neighbor 172.16.255.3 activate neighbor 172.16.255.3 send-community both neighbor 172.16.255.3 route-map BGP-NHU out neighbor 172.16.255.4 activate neighbor 172.16.255.4 send-community both neighbor 172.16.255.4 route-map BGP-NHU out 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 ! route-map BGP-NHU permit 10 set ip next-hop unchanged ! end Spine-02# </pre>

Table 19: Configuring Leaf Switch 1 and Leaf Switch 2 using eBGP when the Spine Switches are in one Autonomous System and each Leaf Switch is in a Different Autonomous System

Leaf Switch 1	Leaf Switch 2
<pre> 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 ! 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 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 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 </pre>	<pre> 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 ! 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 replication-type ingress ! 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 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 ! </pre>

Leaf Switch 1	Leaf Switch 2
<pre> 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 Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 mcast-group 225.0.0.101 member vni 10102 ingress-replication member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.3 ! router bgp 65002 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 ebgp-multihop 255 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 ebgp-multihop 255 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn rewrite-evpn-rt-asn 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 ! </pre>	<pre> 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 Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ! 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 mcast-group 225.0.0.101 member vni 50901 vrf green member vni 10102 ingress-replication ! router ospf 1 router-id 172.16.255.4 ! router bgp 65003 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 65001 neighbor 172.16.255.1 ebgp-multihop 255 neighbor 172.16.255.1 update-source Loopback0 neighbor 172.16.255.2 remote-as 65001 neighbor 172.16.255.2 ebgp-multihop 255 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn rewrite-evpn-rt-asn 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 ! </pre>

Leaf Switch 1	Leaf Switch 2
<pre> 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 ! end Leaf-01# </pre>	<pre> address-family ipv6 vrf green redistribute connected redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! end Leaf-02# </pre>

The following examples provide sample outputs for **show** commands on the devices in the topology configured in the preceding tables:

Spine Switch 1

The following example shows the output for the **show ip ospf neighbor** command on spine switch 1:

```

Spine-01# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address         Interface
172.16.255.4    0     FULL/ -         00:00:34   172.16.14.4    GigabitEthernet1/0/2
172.16.255.3    0     FULL/ -         00:00:38   172.16.13.3    GigabitEthernet1/0/1
    
```

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 19, main routing table version 19
18 network entries using 6192 bytes of memory
38 path entries using 7904 bytes of memory
45/15 BGP path/bestpath attribute entries using 12960 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
24 BGP extended community entries using 1280 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 28384 total bytes of memory
BGP activity 94/76 prefixes, 293/255 paths, scan interval 60 secs
18 networks peaked at 21:10:53 Jun 4 2020 UTC (2d23h ago)

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
172.16.255.3  4      65002   35     27      19    0    0 00:08:54      9
172.16.255.4  4      65003   34     27      19    0    0 00:08:54     11
    
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 1 for route type 2 and the IP address of host device 2:

```

Spine-01# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 10
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    
```

```

11          13
Refresh Epoch 1
65003
 172.16.254.4 (metric 2) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
    Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0
  net: 0x7F7898C7FEF0, path: 0x7F7898C8E578, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on Jun 7 2020 20:42:32 UTC
Refresh Epoch 2
65003
 172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
  Origin incomplete, metric 0, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
    Router MAC:7C21.0DBD.9548
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F7898C7FEF0, path: 0x7F7898C8E728, pathext: 0x7F7898CAE8E0
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on Jun 7 2020 20:41:30 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 1 for route type 2 and the IP address of host device 1:

```

Spine-01# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
  version 4
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    11          13
Refresh Epoch 1
65002
 172.16.254.3 (metric 2) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, metric 0, localpref 100, valid, internal
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0
  net: 0x7F7898C7F290, path: 0x7F7898C8FEC8, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on Jun 7 2020 20:42:32 UTC
Refresh Epoch 2
65002
 172.16.254.3 (metric 2) (via default) from 172.16.255.3 (172.16.255.3)
  Origin incomplete, metric 0, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65001:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F7898C7F290, path: 0x7F7898C8E218, pathext: 0x7F7898CAEE20
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on Jun 7 2020 20:41:30 UTC

```

The following example shows the output for the **show ip pim neighbor** command on spine switch 1:

```

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
172.16.13.3   GigabitEthernet1/0/1   1w4d/00:01:37   v2    1 / DR S P G
172.16.14.4   GigabitEthernet1/0/2   1w4d/00:01:39   v2    1 / DR S P G

```

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 (?)

```

The following example shows the output for the **show ip ro** command on spine switch 1:

```

Spine-01# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Loopback2
    Route metric is 0, traffic share count is 1

```

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

```

Spine-01# 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         1w4d     0      2      ?

```

The following example shows the output for the **show ip msdp sa-cache** command on spine switch 1:

```

Spine-01# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(172.16.254.3, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:01:07/00:05:06, Peer 172.16.254.2
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:00:45/00:05:14, Peer 172.16.254.2

```

The following example shows the output for the **show ip rpf** command on spine switch 1:

```

Spine-01# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: Loopback2
  RPF neighbor: ? (172.16.255.255) - directly connected
  RPF route/mask: 172.16.255.255/32
  RPF type: multicast (connected)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base

```

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

```

Spine-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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 00:01:07/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), 00:00:45/00:02:14, 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), 00:01:07/00:01:52, flags: PA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.13.3
  Outgoing interface list: Null

```

Spine Switch 2

The following example shows the output for the **show ip ospf neighbor** command on spine switch 2:

```

Spine-02# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address        Interface
172.16.255.4    0    FULL/ -         00:00:32   172.16.24.4   GigabitEthernet1/0/2
172.16.255.3    0    FULL/ -         00:00:34   172.16.23.3   GigabitEthernet1/0/1

```

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 19, main routing table version 19
18 network entries using 6192 bytes of memory
38 path entries using 7904 bytes of memory
45/15 BGP path/bestpath attribute entries using 12960 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
24 BGP extended community entries using 1280 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 28384 total bytes of memory
BGP activity 56/38 prefixes, 244/206 paths, scan interval 60 secs
18 networks peaked at 21:11:25 Jun 4 2020 UTC (2d23h ago)

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

```



```

172.16.255.3    4          65002    30       27       19       0       0 00:08:54      9
172.16.255.4    4          65003    30       27       19       0       0 00:08:54     11

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 2 for route type 2 and the IP address of host device 2:

```

Spine-02# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 10
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    9          10
  Refresh Epoch 1
  65003
    172.16.254.4 (metric 2) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      rx pathid: 0, tx pathid: 0
      net: 0x7FB6494C8550, path: 0x7FB64B6D21A8, pathext: 0x0
      flags: net: 0x0, path: 0x3, pathext: 0x0
      Updated on Jun 7 2020 20:43:06 UTC
  Refresh Epoch 2
  65003
    172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      rx pathid: 0, tx pathid: 0x0
      net: 0x7FB6494C8550, path: 0x7FB64B6D3870, pathext: 0x7FB6494D8788
      flags: net: 0x0, path: 0x3, pathext: 0x81
      Updated on Jun 7 2020 20:42:08 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command on spine switch 2 for route type 2 and the IP address of host device 1:

```

Spine-02# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 4
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    9          10
  Refresh Epoch 1
  65002
    172.16.254.3 (metric 2) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
      rx pathid: 0, tx pathid: 0
      net: 0x7FB6494C86B0, path: 0x7FB64B6D25E0, pathext: 0x0
      flags: net: 0x0, path: 0x3, pathext: 0x0
      Updated on Jun 7 2020 20:43:06 UTC
  Refresh Epoch 2
  65002
    172.16.254.3 (metric 2) (via default) from 172.16.255.3 (172.16.255.3)
      Origin incomplete, metric 0, localpref 100, valid, external, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8

```

```

rx pathid: 0, tx pathid: 0x0
net: 0x7FB6494C86B0, path: 0x7FB64B6D31B0, pathext: 0x7FB6494D8CC8
flags: net: 0x0, path: 0x3, pathext: 0x81
Updated on Jun 7 2020 20:42:08 UTC

```

The following example shows the output for the **show ip pim neighbor** command on spine switch 2:

```

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
172.16.23.3       GigabitEthernet1/0/1  00:34:48/00:01:27 v2    1 / DR S P G
172.16.24.4       GigabitEthernet1/0/2  1w4d/00:01:36    v2    1 / DR S P G

```

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 (?)

```

The following example shows the output for the **show ip ro** command on spine switch 2:

```

Spine-02# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "connected", distance 0, metric 0 (connected, via interface)
  Routing Descriptor Blocks:
  * directly connected, via Loopback2
    Route metric is 0, traffic share count is 1

```

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

```

Spine-02# 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       1w4d    0         2       ?

```

The following example shows the output for the **show ip msdp sa-cache** command on spine switch 2:

```

Spine-02# show ip msdp sa-cache
RPF information for ? (172.16.255.255)
  RPF interface: Loopback2
  RPF neighbor: ? (172.16.255.255) - directly connected
  RPF route/mask: 172.16.255.255/32
  RPF type: multicast (connected)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base

```

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

```
Spine-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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 1w4d/00:03:27, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:34:36/00:03:22
    GigabitEthernet1/0/2, Forward/Sparse, 2d23h/00:03:27

(172.16.254.4, 225.0.0.101), 00:00:50/00:02:09, flags: A
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 00:00:50/00:03:22

(172.16.254.3, 225.0.0.101), 00:01:11/00:01:47, flags: A
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:01:11/00:03:27
```

Leaf Switch 1

The following example shows the output for the **show ip ospf neighbor** command on leaf switch 1:

```
Leaf-01# show ip ospf neighbor
Neighbor ID      Pri  State           Dead Time   Address      Interface
172.16.255.2    0    FULL/ -         00:00:31   172.16.23.2  GigabitEthernet1/0/2
172.16.255.1    0    FULL/ -         00:00:34   172.16.13.1  GigabitEthernet1/0/1
```

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

```
Leaf-01# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.3, local AS number 65002
BGP table version is 99, main routing table version 99
27 network entries using 9288 bytes of memory
36 path entries using 7488 bytes of memory
22/15 BGP path/bestpath attribute entries using 6336 bytes of memory
1 BGP AS-PATH entries using 40 bytes of memory
18 BGP extended community entries using 944 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 24096 total bytes of memory
BGP activity 483/450 prefixes, 1123/1081 paths, scan interval 60 secs
```

27 networks peaked at 13:15:47 May 26 2020 UTC (1w5d ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	27	34	99	0	0	00:08:30	9
172.16.255.2	4	65001	27	29	99	0	0	00:08:25	9

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 host device 2:

```
Leaf-01# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 93
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  65001 65003, imported path from
[2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24 (global)
  172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
    Origin incomplete, localpref 100, valid, external, best
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
    rx pathid: 0, tx pathid: 0x0
    net: 0x7F575E4795D0, path: 0x7F575E1EFC38, pathext: 0x7F575E201308, exp_net:
0x7F575E47AA70
    flags: net: 0x0, path: 0x40000000000003, pathext: 0x81
    Updated on Jun 7 2020 20:40:17 UTC
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 84
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    21
  Refresh Epoch 1
  65001 65003
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
    Origin incomplete, localpref 100, valid, external
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
```

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 host device 1:

```
Leaf-01# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 4
Paths: (1 available, best #1, table evi_101)
  Advertised to update-groups:
    21
  Refresh Epoch 1
  Local
  :: (via default) from 0.0.0.0 (172.16.255.3)
    Origin incomplete, localpref 100, weight 32768, valid, sourced, local, best
    EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
    Extended Community: RT:1:1 RT:65002:101 ENCAP:8
      Router MAC:10B3.D56A.8FC8
    Local irb vxlan vtep:
      vrf:green, l3-vni:50901
      local router mac:10B3.D56A.8FC8
      core-irb interface:Vlan901
      vtep-ip:172.16.254.3
    rx pathid: 0, tx pathid: 0x0
    net: 0x7F575E47ABD0, path: 0x7F575E1F13D8, pathext: 0x7F575E201968
```

```
flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
Updated on Jun 4 2020 21:26:02 UTC
```

The following example shows the output for the **show l2vpn evpn mac ip** command on leaf switch 1:

```
Leaf-01# show l2vpn evpn mac ip
IP Address          EVI   VLAN  MAC Address      Next Hop(s)
-----
10.1.101.11        101   101   f4cf.e243.34c1  Gi1/0/10:101
10.1.101.12        101   101   44d3.ca28.6cc1  172.16.254.4
10.1.102.11        102   102   f4cf.e243.34c2  Gi1/0/10:102
10.1.102.12        102   102   44d3.ca28.6cc2  172.16.254.4
```

The following example shows the output for the **show ip pim neighbor** command on leaf switch 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
172.16.13.1   GigabitEthernet1/0/1  1w4d/00:01:17    v2   1 / S P G
172.16.23.2   GigabitEthernet1/0/2  00:34:19/00:01:24 v2   1 / S P G
```

The following example shows the output for the **show ip pim rp mapping** command on leaf switch 1:

```
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 (?)
```

The following example shows the output for the **show ip ro** command on leaf switch 1:

```
Leaf-01# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 172.16.23.2 on GigabitEthernet1/0/2, 00:34:08 ago
  Routing Descriptor Blocks:
    172.16.23.2, from 172.16.255.2, 00:34:08 ago, via GigabitEthernet1/0/2
      Route metric is 2, traffic share count is 1
    * 172.16.13.1, from 172.16.255.1, 1w4d ago, via GigabitEthernet1/0/1
      Route metric is 2, traffic share count is 1
```

The following example shows the output for the **show ip rpf** command on leaf switch 1:

```
Leaf-01# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
  RPF interface: GigabitEthernet1/0/2
  RPF neighbor: ? (172.16.23.2)
  RPF route/mask: 172.16.255.255/32
  RPF type: unicast (ospf 1)
  Doing distance-preferred lookups across tables
  RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

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

```
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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 8w2d/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, 1w5d/00:01:01

(172.16.254.4, 225.0.0.101), 00:00:21/00:02:38, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 00:00:21/00:02:38

(172.16.254.3, 225.0.0.101), 00:00:43/00:02:46, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 00:00:43/00:02:46
```

Leaf Switch 2

The following example shows the output for the **show ip ospf neighbor** command on leaf switch 2:

```
Leaf-02# show ip ospf neighbor
Neighbor ID      Pri   State           Dead Time   Address          Interface
172.16.255.2     0    FULL/ -         00:00:36   172.16.24.2     GigabitEthernet1/0/2
172.16.255.1     0    FULL/ -         00:00:31   172.16.14.1     GigabitEthernet1/0/1
```

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

```
Leaf-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.4, local AS number 65003
BGP table version is 83, main routing table version 83
25 network entries using 8600 bytes of memory
36 path entries using 7488 bytes of memory
23/15 BGP path/bestpath attribute entries using 6624 bytes of memory
1 BGP AS-PATH entries using 40 bytes of memory
19 BGP extended community entries using 984 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 23736 total bytes of memory
BGP activity 95/64 prefixes, 207/163 paths, scan interval 60 secs
25 networks peaked at 21:31:21 Jun 4 2020 UTC (2d23h ago)
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	27	34	83	0	0	00:08:40	9
172.16.255.2	4	65001	27	29	83	0	0	00:08:35	9

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 host device 2:

```
Leaf-02# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.12
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.12]/24,
version 4
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
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65003:101 ENCAP:8
  Router MAC:7C21.0DBD.9548
  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
  net: 0x7F84B8F2D778, path: 0x7F84BB3149F0, pathext: 0x7F84BB526788
  flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
  Updated on Jun 4 2020 21:30:20 UTC
```

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 host device 1:

```
Leaf-02# show bgp l2vpn evpn route-type 2 0 f4cfe24334c1 10.1.101.11
BGP routing table entry for [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 72
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Advertised to update-groups:
    2
  Refresh Epoch 1
  65001 65002
  172.16.254.3 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
  Origin incomplete, localpref 100, valid, external
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65003:101 ENCAP:8
  Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0
  net: 0x7F84B8F2E958, path: 0x7F84BB313FD0, pathext: 0x0
  flags: net: 0x0, path: 0x3, pathext: 0x0
  Updated on Jun 7 2020 20:44:45 UTC
  Refresh Epoch 1
  65001 65002
  172.16.254.3 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65003:101 ENCAP:8
  Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F84B8F2E958, path: 0x7F84BB313178, pathext: 0x7F84BB526548
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on Jun 7 2020 20:44:44 UTC
```

```

BGP routing table entry for [2][172.16.254.4:101][0][48][F4CFE24334C1][32][10.1.101.11]/24,
version 78
Paths: (1 available, best #1, table evi_101)
  Not advertised to any peer
  Refresh Epoch 1
  65001 65002, imported path from
[2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24 (global)
  172.16.254.3 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
  Origin incomplete, localpref 100, valid, external, best
  EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
  Extended Community: RT:1:1 RT:65003:101 ENCAP:8
    Router MAC:10B3.D56A.8FC8
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F84B8F2D358, path: 0x7F84BB314258, pathext: 0x7F84BB5265A8, exp_net:
0x7F84B8F2E958
  flags: net: 0x0, path: 0x40000000000003, pathext: 0x81
  Updated on Jun 7 2020 20:44:44 UTC

```

The following example shows the output for the **show l2vpn evpn mac ip** command on leaf switch 2:

```

Leaf-02# show l2vpn evpn mac ip
IP Address          EVI    VLAN  MAC Address      Next Hop(s)
-----
10.1.101.11         101    101   f4cf.e243.34c1   172.16.254.3
10.1.101.12         101    101   44d3.ca28.6cc1   Gi1/0/10:101
10.1.102.11         102    102   f4cf.e243.34c2   172.16.254.3
10.1.102.12         102    102   44d3.ca28.6cc2   Gi1/0/10:102

```

The following example shows the output for the **show ip pim neighbor** command on leaf switch 2:

```

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
172.16.14.1       GigabitEthernet1/0/1  1w4d/00:01:42    v2   1 / S P G
172.16.24.2       GigabitEthernet1/0/2  1w4d/00:01:19    v2   1 / S P G

```

The following example shows the output for the **show ip pim rp mapping** command on leaf switch 2:

```

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 (?)

```

The following example shows the output for the **show ip ro** command on leaf switch 2:

```

Leaf-02# show ip ro 172.16.255.255
Routing entry for 172.16.255.255/32
  Known via "ospf 1", distance 110, metric 2, type intra area
  Last update from 172.16.14.1 on GigabitEthernet1/0/1, 3d00h ago
  Routing Descriptor Blocks:
  * 172.16.24.2, from 172.16.255.2, 3d00h ago, via GigabitEthernet1/0/2
    Route metric is 2, traffic share count is 1

```



```
172.16.14.1, from 172.16.255.1, 3d00h ago, via GigabitEthernet1/0/1
Route metric is 2, traffic share count is 1
```

The following example shows the output for the **show ip rpf** command on leaf switch 2:

```
Leaf-02# show ip rpf 172.16.255.255
RPF information for ? (172.16.255.255)
RPF interface: GigabitEthernet1/0/2
RPF neighbor: ? (172.16.24.2)
RPF route/mask: 172.16.255.255/32
RPF type: unicast (ospf 1)
Doing distance-preferred lookups across tables
RPF topology: ipv4 multicast base, originated from ipv4 unicast base
```

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

```
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
Outgoing interface flags: H - Hardware switched, A - Assert winner, p - PIM Join
Timers: Uptime/Expires
Interface state: Interface, Next-Hop or VCD, State/Mode

(*, 225.0.0.101), 2w3d/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, 2d23h/00:01:40

(172.16.254.4, 225.0.0.101), 00:00:31/00:02:58, flags: FTx
Incoming interface: Loopback1, RPF nbr 0.0.0.0
Outgoing interface list:
GigabitEthernet1/0/2, Forward/Sparse, 00:00:31/00:02:58

(172.16.254.3, 225.0.0.101), 00:00:52/00:02:07, flags: JTx
Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
Outgoing interface list:
Tunnel0, Forward/Sparse-Dense, 00:00:52/00:02:07
```




CHAPTER 6

Configuring DHCP Relay in a BGP EVPN VXLAN Fabric

- [Information About DHCP Relay in a BGP EVPN VXLAN Fabric, on page 197](#)
- [How to Configure DHCP Relay in a BGP EVPN VXLAN Fabric, on page 199](#)
- [Configuration Examples for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 203](#)
- [Additional References for DHCP Relay in a BGP EVPN VXLAN Fabric, on page 221](#)

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 DHCP 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

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 204](#) 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 209](#) 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 212](#) 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 217](#) for a configuration example.

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 [How to Configure 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 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:	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].

	Command or Action	Purpose
	<pre>Device(config)# ip dhcp compatibility suboption link-selection standard Device(config)# ip dhcp compatibility suboption server-override standard</pre>	Use the server-override standard keyword to switch to standard DHCP option 82[11].
Step 7	<pre>ip dhcp snooping vlan <i>vlan-id-list</i> Example: Device(config)# ip dhcp snooping vlan 201-202</pre>	Enables DHCP snooping on the specified list of VLANs.
Step 8	<pre>ip dhcp snooping Example: Device(config)# ip dhcp snooping</pre>	Enables DHCP snooping globally on the VTEP.
Step 9	<pre>end Example: Device(config)# end</pre>	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	<pre>enable Example: Device> enable</pre>	Enables privileged EXEC mode. Enter your password, if prompted.
Step 2	<pre>configure terminal Example: Device# configure terminal</pre>	Enters global configuration mode.
Step 3	<pre>interface vlan <i>vlan-id</i> Example: Device(config)# interface Vlan 201</pre>	Enters interface configuration mode for the specified VLAN interface. This VLAN interface acts as the GiAddr.
Step 4	<pre>vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding green</pre>	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.

	Command or Action	Purpose
Step 5	ip dhcp relay information option vpn-id Example: <pre>Device(config-if)# ip dhcp relay information option vpn-id</pre>	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: <pre>Device(config-if)# ip dhcp relay source-interface Loopback13</pre>	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: <pre>Device(config-if)# ip address 192.168.1.201 255.255.255.0</pre>	Sets the IP address for the VLAN interface.
Step 8	ip helper-address [global vrf vrf-name] ip-address Example: <pre>Device(config-if)# ip helper-address global 192.168.3.100 Device(config-if)# ip helper-address vrf green 192.168.20.20</pre>	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 reachable over the tenant VRF.
Step 9	exit Example: <pre>Device(config-if)# exit</pre>	Exits interface configuration mode and returns to global configuration mode.
Step 10	end Example: <pre>Device(config)# end</pre>	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 <i>vlan-id</i> Example: Device(config)# interface vlan 203	Enters interface configuration mode for the specified VLAN interface.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding green	Configures the SVI for the VLAN and associates the specified VRF with the interface.
Step 5	ip address <i>ip-address</i> 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 <i>ipv6-address</i> 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 <i>interface-id</i> Example: Device(config)# interface GigabitEthernet1/0/30	Enters interface configuration mode for the specified interface.
Step 10	switchport access vlan <i>vlan-id</i> Example:	Specifies the VLAN to be used as access VLAN when the interface is in access mode.

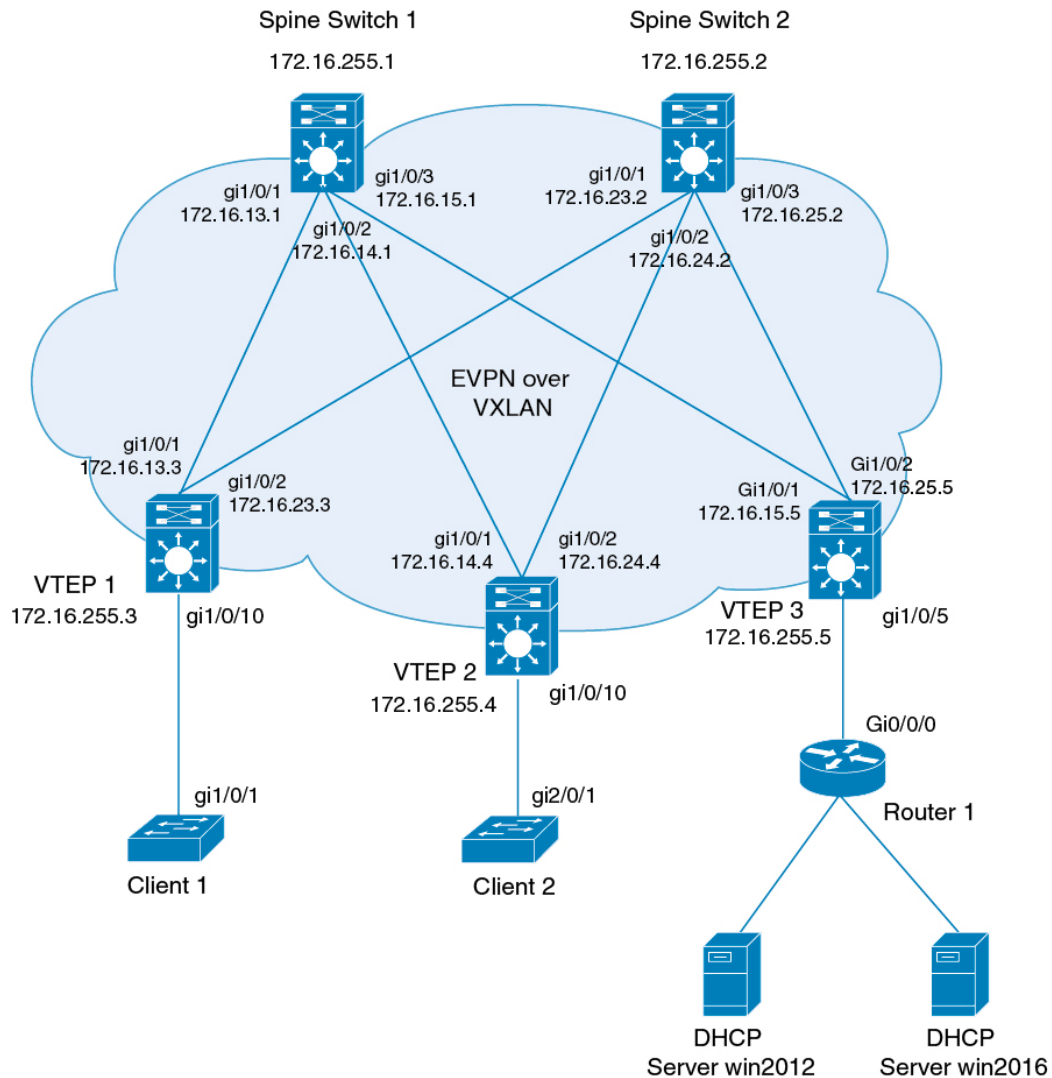
	Command or Action	Purpose
	Device(config-if)# switchport access vlan 203	
Step 11	switchport mode access Example: Device(config-if)# switchport mode access	Configures the interface as an access interface.
Step 12	exit Example: Device(config-if)# exit	Exits interface configuration mode and returns to 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 [Figure 18: DHCP Relay Deployment in a BGP EVPN VXLAN Fabric](#), on [page 204](#).

- Example: DHCP Server is in the Layer 3 Default VRF and the DHCP Client is in the Tenant VRF
- Example: DHCP Server and DHCP Client are in the Same Tenant VRF
- Example: DHCP Client and DHCP Server are in Different Tenant VRFs
- Example: DHCP Server is in a non-Default, non-VXLAN VRF and DHCP Client is in the Tenant VRF

Figure 18: 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 [Figure 18: DHCP Relay Deployment in a BGP EVPN VXLAN Fabric](#) 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:

Table 20: 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 21: Configuring VTEP 1 when DHCP Server is in the Layer 3 Default VRF and DHCP Client is in the Tenant VRF

VTEP 1
<pre> 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 ! </pre>

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 203](#).

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 [Figure 18: DHCP Relay Deployment in a BGP EVPN VXLAN Fabric](#) 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:

Table 22: Configuring DHCP when DHCP Server and DHCP Client are in the Same Tenant VRF

DHCP Configuration Snippet
<pre><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></pre>

Table 23: Configuring VTEP 1 when DHCP Server and DHCP Client are in the Same Tenant VRF

VTEP 1
<pre> 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 ! </pre>

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 203.

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 [Figure 18: DHCP Relay Deployment in a BGP EVPN VXLAN Fabric](#) 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:

Table 24: 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 25: Configuring VTEP 1 when DHCP Server and DHCP Client are in Different Tenant VRFs

VTEP 1
<pre> 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 ! </pre>

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 203](#).

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 [Figure 18: DHCP Relay Deployment in a BGP EVPN VXLAN Fabric](#) 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:

Table 26: 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 27: 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
<pre> 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 </pre>

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 203.

Additional References for DHCP Relay in a BGP EVPN VXLAN Fabric

Related Documents

Related Topic	Document Title
DHCP Server configuration in IOS XE EVPN VXLAN	Configure DHCP in IOS XE EVPN/VXLAN



CHAPTER

7

Configuring VXLAN-Aware Flexible Netflow

- [Restrictions for VXLAN-Aware Flexible NetFlow, on page 223](#)
- [Information About VXLAN-Aware Flexible NetFlow, on page 223](#)
- [How to Configure VXLAN-Aware Flexible NetFlow, on page 223](#)
- [Configuration Examples for VXLAN-Aware Flexible NetFlow, on page 228](#)

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 <i>flow-record-name</i> 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 <i>flow-exporter-name</i> Example: Device(config)# flow exporter e1	Creates a flow exporter and enters flow exporter configuration mode.
Step 4	destination <i>ipv4-address</i> Example: Device(config-flow-exporter)# destination 172.16.103.2	Sets the IPv4 destination address or hostname for the exporter.
Step 5	source <i>interface-type interface-number</i> 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 <i>seconds</i> 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 <i>port-number</i> 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 <i>flow-monitor-name</i> 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 <i>flow-exporter-name</i> 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 <i>flow-record-name</i> 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# <code>configure terminal</code>	Enters global configuration mode.
Step 3	interface <i>nve-interface-number</i> Example: Device(config)# <code>interface nve1</code>	Specifies the network virtualization edge (NVE) interface number and enters interface configuration mode.
Step 4	ip flow monitor <i>flow-monitor-name</i> {input output} Example: Device(config-if)# <code>ip flow monitor vxlan_nf_monitor_input input</code>	Associates the IPv4 flow monitor to the NVE interface for input or output packets.
Step 5	ipv6 flow monitor <i>flow-monitor-name</i> {input output} Example: Device(config-if)# <code>ipv6 flow monitor vxlan_nf_v6monitor_input input</code>	Associates the IPv6 flow monitor to the NVE interface for input or output packets.
Step 6	end Example: Device(config-if)# <code>end</code>	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



357052

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 28: Configuring VTEP 1 to enable VXLAN-Aware Flexible NetFlow

VTEP 1
<pre> 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 ! </pre>

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
!
!
```



```

12 AAAA.CCCC.1004 AAAA.BBBB.1004 3
(13vni5001) 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

```




CHAPTER 8

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 235](#)
- [How to Configure BGP EVPN VXLAN QoS, on page 236](#)
- [BGP EVPN VXLAN QoS Modes, on page 237](#)
- [BGP EVPN VXLAN Traffic without QoS Policy , on page 238](#)
- [BGP EVPN VXLAN QoS Marking with an Input Policy, on page 239](#)
- [BGP EVPN VXLAN QoS Marking with Input and Output Policies, on page 240](#)

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	<p>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.</p> <p>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.</p>
Shaping and Policing	<p>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.</p>
Queuing	<p>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.</p>
Bandwidth	<p>Bandwidth allocation determines the available capacity for traffic that is subject to QoS policies.</p>
Trust	<p>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.</p>
Class of Service (CoS)	<p>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.</p>
Differentiated Services Code Point (DSCP)	<p>The first six bits of the ToS byte in the IP header. DSCP is only present in an IP packet.</p>
Explicit Congestion Notification (ECN)	<p>The last two bits of the ToS byte in the IP header. ECN is only present in an IP packet.</p>

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 9

Configuring Optimized Layer 2 Overlay Multicast

- [Restrictions for Optimized Layer 2 Overlay Multicast, on page 241](#)
- [Information About Optimized Layer 2 Overlay Multicast, on page 241](#)
- [How to Configure Optimized Layer 2 Overlay Multicast, on page 245](#)
- [Verifying Optimized Layer 2 Overlay Multicast on VTEP and Centralized Gateway, on page 249](#)
- [Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 250](#)

Restrictions for Optimized Layer 2 Overlay Multicast

- This feature does not support data MDT in the underlay.
- Optimized Layer 2 overlay multicast source specific 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 Centralized Gateway that extends the bridge-domain over 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 IGMP or MLD Proxy for EVPN, as defined in [draft-ietf-bess-evpn-igmp-ml-d-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 (EVI) 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 EVI, 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 to show 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.

487154

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 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, both could be on the same VTEP or they could be on different VTEPs within the EVPN VXLAN fabric.

Source and Receiver on Different VLANs, connected to Different Layer 2 VTEPs

Figure 24: Topology to show Layer 2 Overlay Multicast Traffic between Different VLANs



467152

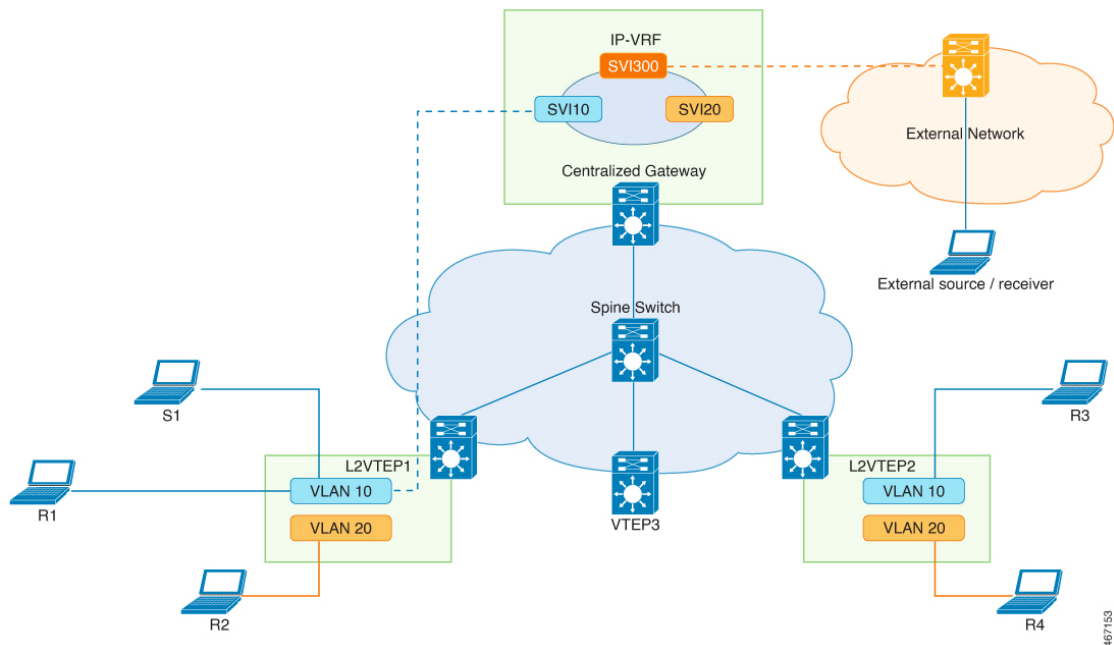
1. CGW sends EVPN route type 6 (*, *) to all the Layer 2 VTEPs in EVPN VXLAN fabric.
2. Receiver R4 in VLAN 20 sends (*, G) IGMP join to VTEP 2. (*, G) is created in IGMP snooping table on VTEP 2.
3. VTEP 2 sends EVPN route type 6 (*, G) to the CGW and all other Layer 2 VTEPs. (*, G) is also created 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 to show 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 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 multdestination Layer 2 overlay broadcast, unknown unicast, and multicast (BUM) traffic. For more information, see [Ingress Replication, on page 14](#) in the “Configuring EVPN VXLAN Layer 2 Overlay Network” chapter. 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 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](#) in the “Configuring EVPN VXLAN Layer 2 Overlay Network” chapter. 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

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: Device# configure terminal	Enters global configuration mode.
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 [Prerequisites for Configuring Optimized Layer 2 Overlay Multicast](#), 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.

	Command or Action	Purpose
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 <i>version-num</i> Example: Device(config)# ip igmp snooping querier version 3	Configures the IGMP version which the querier feature uses. By default, IGMP version 2 is configured. IGMP version 2 extends IGMP functionality with features such as the IGMP leave process, group-specific queries, and an explicit maximum query response time. IGMP version 3 provides 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 <i>seconds</i>] [timer expiry <i>seconds</i>] 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 1 to 25 seconds. timer expiry sets the length of time until the IGMP querier expires. The range is 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 [Prerequisites for Configuring Optimized Layer 2 Overlay Multicast](#), 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.

	Command or Action	Purpose
Step 3	ipv6 mld snooping Example: Device(config)# ipv6 mld snooping	Enables the MLD snooping.
Step 4	ipv6 mld snooping querier Example: Device(config)# ipv6 mld snooping querier	Enables the 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

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 108](#).
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	Enables IP multicast routing.

	Command or Action	Purpose
	Example: Device(config)# ip multicast-routing	Note If the Centralized Gateway is associated to a VRF, run the ip multicast-routing vrf vrf-name command.
Step 4	interface vlan <i>vlan-id</i> 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 <i>version</i> 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 Steps 4–6 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 <i>vrf-name</i>] rp-address <i>rp-address</i> 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 IGMP Snooping querier version is as expected.
show ip igmp snooping groups	Verify the local and remote multicast group membership in IGMP Snooping.

Command	Purpose
<code>show ip mroute vrf <vrf></code>	Verify the MVPN routing and forwarding instance for the specified VRF.
<code>show ip igmp vrf <vrf-name> groups detail</code>	Verify the multicast groups with receivers that were learned through IGMP, on the centralized gateway.
<code>show ipv6 mld snooping querier</code>	Verify that MLD Snooping querier version is as expected.
<code>show ipv6 mld snooping membership</code>	Verify the local and remote multicast group membership in MLD Snooping.
<code>show ipv6 mroute vrf <vrf-name> <multicast-group-address></code>	Verify the forwarding status of each IPv6 multicast route.
<code>show ipv6 mld vrf <vrf-name> groups detail</code>	Verify the multicast groups with receivers that were learned through MLD, on the centralized gateway.
<code>show l2vpn evpn default-gateway</code>	Verify the entries of the default gateway database.
<code>show l2vpn evpn multicast [local] [remote]</code>	Verify the local and remote multicast group membership in EVPN Manager.
<code>show l2vpn evpn evi <evpn-instance> detail</code>	Verify Layer 2 multicast status for the specified EVPN instance.
<code>show l2route evpn multicast smet [local] [remote] [local combined]</code>	Verify the SMET routes that the Layer 2 Route Information Base (L2RIB) receives from EVPN.
<code>show l2route evpn multicast routes group <multicast-address></code>	Verify the multicast routes that are sent to L2FIB.
<code>show bgp l2vpn evpn route-type 6</code>	Verify the local and remote SMET routes in BGP.
<code>show bgp l2vpn evpn summary</code>	Verify the BGP EVPN routes on the Spine switch.
<code>show ip pim neighbor</code>	Verify the Protocol Independent Multicast (PIM) neighbors that are discovered by the device.
<code>show ip pim rp mapping</code>	Verify that the PIM Group-to-Rendezvous Point (RP) mappings are populated correctly on the device.
<code>show ip msdp peer</code>	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:

Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication

Table 29: 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
<pre>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</pre>	<pre>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</pre>	<pre>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 TenGigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface</pre>

Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication

VTEP 1	CGW	VTEP 3
<pre> ! 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 </pre>	<pre> ! 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 ! 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 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 </pre>	<pre> TenGigabitEthernet1/0/11 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 </pre>

VTEP 1	CGW	VTEP 3
	<pre>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</pre>	

Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication

Table 30: 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
<pre> Spine-01#show running-config hostname Spine-01 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 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 172.16.255.2 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 </pre>	<pre> Spine-02#show running-config hostname Spine-02 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.23.2 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.2 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 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-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 </pre>

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

Return to [Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication, on page 251](#)

Return to [Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 250](#).

Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay

The following sections provide sample outputs for **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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay](#), on page 258

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	V1102

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	Gi1/0/11, 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, 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:

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay](#), on page 258

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	Tel1/0/11, 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:
  TenGigabitEthernet1/0/11 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	Tel1/0/11	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	*	Tel1/0/11:102, 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 Te1/0/11:102 239.1.1.2 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: 172.16.255.4, 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: 172.16.255.4, 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: 172.16.255.4, 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
```


Return to [Example 1: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Ingress Replication](#)

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 [Figure 26: Topology for Optimized Layer 2 Overlay Multicast for IPv4 Traffic with Ingress Replication](#).

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 31: Configure VTEP 1, CGW, and VTEP 3 for Optimized Layer 2 Multicast for IPv4 and IPv6, with Ingress Replication

VTEP1	CGW	VTEP 3
-------	-----	--------

Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6

VTEP1	CGW	VTEP 3
<pre>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 !</pre>	<pre>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 !</pre>	<pre>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 TenGigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point</pre>

VTEP1	CGW	VTEP 3
<pre> 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 </pre>	<pre> 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 </pre>	<pre> ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface TenGigabitEthernet1/0/11 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 </pre>

Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6

VTEP1	CGW	VTEP 3
	<pre> ! 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 </pre>	

Table 32: 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
----------------	----------------

Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6

Spine Switch 1	Spine Switch 2
<pre> Spine-01#show running-config hostname Spine-01 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 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 172.16.255.2 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 </pre>	<pre> Spine-02#show running-config hostname Spine-02 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.23.2 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.2 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 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-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 </pre>

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

Return to [Example 2: Configuring Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6](#), on page 266.

Return to [Configuration Examples for Optimized Layer 2 Overlay Multicast](#), on page 250.

Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic

The following sections provide sample outputs for **show** commands to verify Optimized Layer 2 Overlay Multicast with Ingress Replication on the devices in the topology configured above.

To see the outputs of **show** commands for IPv4 multicast traffic, refer [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication in the Underlay](#), on page 258

The outputs of **show** commands for IPv6 multicast traffic are the following:

[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

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

```

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::239:1:1:2**

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 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: 172.16.255.4, 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
```

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 273](#)

Outputs to verify configuration on CGW

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  V1101
Y   2001:10:1:101::1                    101   101    7c21.0dbd.9541  V1101
```

```
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 Last-Join/ -----	Interface	Reporter	Vlan	Uptime Last-Leave
::/FF06::239:1:1:1 00:00:05 /	Gi1/0/10	FE80::46D3:CAFF:FE28:6CC1	101	00:00:00 00:00:05
::/FF06:239:1:1::1 20:16:44 /	Gi1/0/10	FE80::46D3:CAFF:FE28:6CC1	101	00:00:47 20:15:56
::/FF06::239:1:1:1 20:15:12 /	Tu0	FE80::AC10:FE06	101	20:15:12 -
::/FF06:239:1:1::1 20:15:21 /	Tu0	FE80::AC10:FE06	101	00:00:00 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

```

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  V1102

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-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][172.16.254.6:102][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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 273](#)

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:BDFE: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 Last-Join/	Interface Reporter	Vlan Uptime Last-Leave
::/FF06::239:1:1:1 00:00:57 /	Tel/0/10 FE80::EEE1:A9FF:FE37:92C1	101 00:00:00 00:00:57
::/FF06:239:1:1::1 20:25:57 /	Tel/0/10 FE80::EEE1:A9FF:FE37:92C1	101 00:01:07 20:24:49
::/FF06::239:1:1:1 20:25:16 /	Tu0 FE80::AC10:FE04	101 20:25:16 -
::/FF06:239:1:1::1 20:26:13 /	Tu0 FE80::AC10:FE04	101 00:00:48 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:
  TenGigabitEthernet1/0/10 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    Te1/0/10           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, Te1/0/10:101

```

```

Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:1
EVI  ETAG    Origin              Group              Filter
Source(s)
-----
101  0       Te1/0/10:101       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

```

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:BDFE: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 \
                                Te1/0/11 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:
  TenGigabitEthernet1/0/11 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	Te1/0/11	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, Te1/0/11:102
```

```
Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:2
```

```
-----
EVI ETAG Origin Group Filter
Source(s)
-----
102 0 Te1/0/11:102 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: 172.16.255.4, 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: 172.16.255.4, 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: 172.16.255.4, 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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 273](#)

Outputs to verify configuration on Spine Switch 1

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

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	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

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

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 273](#)

Outputs to verify configuration on Spine Switch 2

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 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 (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

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

Return to [Verifying Optimized Layer 2 Overlay Multicast with Ingress Replication for IPv4 and IPv6 Traffic, on page 273](#)

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:

Example 3: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication

Table 33: 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
<pre>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</pre>	<pre>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 ! 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</pre>	<pre>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 TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2 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</pre>

Example 3: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication

VTEP 1	CGW	VTEP 3
<pre> 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 </pre>	<pre> 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 </pre>	<pre> ! interface TenGigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface TenGigabitEthernet1/0/11 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 </pre>

VTEP 1	CGW	VTEP 3
	<pre> 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 </pre>	

Example 3: Configuring Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication

Table 34: 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
<pre> Spine-01#show running-config hostname Spine-01 ! ip routing ! ip multicast-routing ! 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 172.16.255.2 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate </pre>	<pre> Spine-02#show running-config hostname Spine-02 ! ip routing ! ip multicast-routing ! 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 172.16.255.1 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate </pre>

Spine Switch 1	Spine Switch 2
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
neighbor 172.16.255.6 route-reflector-client	neighbor 172.16.255.6 route-reflector-client
exit-address-family	exit-address-family
!	!
ip pim rp-address 172.16.255.255	ip pim rp-address 172.16.255.255
ip msdp peer 172.16.254.2 connect-source	ip msdp peer 172.16.254.1 connect-source
Loopback1 remote-as 65001	Loopback1 remote-as 65001
ip msdp cache-sa-state	ip msdp cache-sa-state
!	!
end	end

Return to [Configuration Example for Optimized Layer 2 Overlay Multicast for IPv4 with Underlay Multicast Replication](#)

Return to [Configuration Examples for Optimized Layer 2 Overlay Multicast, on page 250.](#)

Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay

The following sections provide sample outputs for **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
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 12vpn 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 12vpn 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 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)
-----

```

```

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

```

Return to [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
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,
       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  V1101
```

```
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 239.1.1.1
```

EVI	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:
```

Return to [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
172.16.16.1   TenGigabitEthernet1/0/1  00:12:30/00:01:34 v2   1 / S P G
172.16.26.2   TenGigabitEthernet1/0/2  00:12:29/00:01:33 v2   1 / S P G

Leaf-03#show ip pim rp mapping
PIM Group-to-RP Mappings
```

Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay

```
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 - encaps-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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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	Tel0/10, 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:
  TenGigabitEthernet1/0/10 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    Te1/0/10       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      *          Te1/0/10:101, 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      Te1/0/10:101   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)

```

Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication

```

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

```

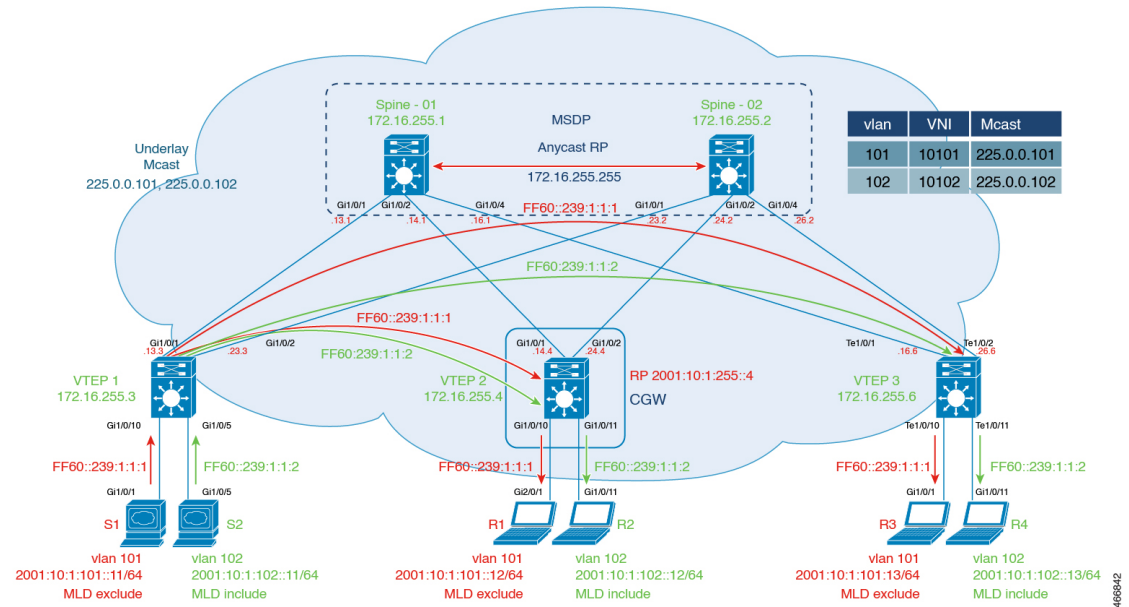
Return to [Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay](#) .

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 [Figure 28: Topology for Optimized Layer 2 Overlay Multicast for IPv4 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:

Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication

Table 35: 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
<pre>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</pre>	<pre>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</pre>	<pre>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 TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2</pre>

Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication

VTEP 1	CGW	VTEP 3
<pre> 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 </pre>	<pre> 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 </pre>	<pre> 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 TenGigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface TenGigabitEthernet1/0/11 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 </pre>

VTEP 1	CGW	VTEP 3
	<pre> 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 </pre>	

Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication

Table 36: 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
<pre> Spine-01#show running-config hostname Spine-01 ! ip routing ! ip multicast-routing ! 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 172.16.255.2 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate </pre>	<pre> Spine-02#show running-config hostname Spine-02 ! ip routing ! ip multicast-routing ! 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 172.16.255.1 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.1 activate </pre>

Spine Switch 1	Spine Switch 2
<pre>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</pre>	<pre>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</pre>

Return to [Example 4: Configuring Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 with Underlay Multicast Replication](#).

Return to [Configuration Examples for Optimized Layer 2 Overlay Multicast](#), on page 250

Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic

The following sections provide sample outputs for **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 outputs of **show** commands for IPv4 multicast traffic, refer [Verifying Optimized Layer 2 Overlay Multicast with Default MDT in the Underlay](#), on page 300.

The outputs of **show** commands for IPv6 multicast traffic are the following:

[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

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
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 Last-Join/	Interface	Reporter	Vlan	Uptime	Last-Leave
::/FF06::239:1:1:1 1d00h /	Tu0	FE80::AC10:FE04	101	1d00h	-
::/FF06::239:1:1:1 1d00h /	Tu0	FE80::AC10:FE06	101	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 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)
-----

```

```

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
```

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
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
Pseudopoports:
  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 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)

```

```

-----
-----

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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 318](#)

Outputs to verify configuration on CGW

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
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  V1101
Y   2001:10:1:101::1                    101   101   7c21.0dbd.9541  V1101

```

```

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/
-----
::/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

```

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 Address	Interface	Uptime/Expires	Ver	DR 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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast with Default MDT for IPv4 and IPv6 Traffic, on page 318](#)

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
172.16.16.1       TenGigabitEthernet1/0/1  1d00h/00:01:29    v2    1 / S P G
172.16.26.2       TenGigabitEthernet1/0/2  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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:03:20, flags:

```

```

(172.16.254.3, 225.0.0.101), 1d00h/00:01:34, flags: JTx
  Incoming interface: TenGigabitEthernet1/0/2, 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 Last-Join/	Interface Reporter	Vlan Uptime	Last-Leave
::/FF06::239:1:1:1	Tel/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:
  TenGigabitEthernet1/0/10 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	Tel/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
101	0	FF06::239:1:1:1	::

Te1/0/10:101, 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           Te1/0/10:101    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 Address	Interface	Uptime/Expires	Ver	DR Prio/Mode
172.16.16.1	TenGigabitEthernet1/0/1	1d00h/00:01:19	v2	1 / S P G
172.16.26.2	TenGigabitEthernet1/0/2	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 - encaps-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: TenGigabitEthernet1/0/2, 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:
```

```
    TenGigabitEthernet1/0/2, Forward/Sparse, 1d00h/00:02:51, flags: A
```

```
(172.16.254.3, 225.0.0.102), 1d00h/00:01:08, flags: JTx
```

```
  Incoming interface: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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 102**

```
IP address      : FE80:0:356:0:E75:BDF: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-Leave
2001:10:1:102::11/FF06::239:1:1:2 \	Te1/0/11	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:
  TenGigabitEthernet1/0/11 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    Te1/0/11        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  ::
      Tel/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        Tel/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 318

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 Address	Interface	Uptime/Expires	Ver	DR 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 - encaps-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 318

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 Address	Interface	Uptime/Expires	Ver	DR 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 - encaps-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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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: GigabitEthernet1/0/2, 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 Address	Interface	Uptime/Expires	Ver	DR 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 - LISF 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: GigabitEthernet1/0/2, 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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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 318

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 handed off to an external network.

Figure 30: Topology for 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. The source is located in Layer 2 VTEP1 and receivers are located in Layer 2 VTEP 3, the Centralized Gateway, and also located in the external network. Centralized gateway is configured on VTEP 2. VLAN 101 has IGMPv2 and VLAN 102 has IGMPv3 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
-------	-----	-------

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

VTEP1	CGW	VTEP3
<pre> 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 </pre>	<pre> 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 GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.182 255.255.255.0 negotiation auto ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.14.4 255.255.255.0 </pre>	<pre> 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 TenGigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! </pre>

VTEP1	CGW	VTEP3
<pre> 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 </pre>	<pre> 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 ! 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 </pre>	<pre> interface TenGigabitEthernet1/0/11 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 </pre>

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

VTEP1	CGW	VTEP3
	<pre> 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 ! end </pre>	

Table 38: Configure Spine Switch 1, Spine Switch 2, and External Device

Spine Switch 1	Spine Switch 2	External Device
		<pre> 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 </pre>

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

Spine Switch 1	Spine Switch 2	External Device
<pre> Spine-01#show running-config hostname Spine-01 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 </pre>	<pre> Spine-02#show running-config hostname Spine-02 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.23.2 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.2 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 </pre>	

Spine Switch 1	Spine Switch 2	External Device
<pre> 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 </pre>	<pre> 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 </pre>	

Return to [Example 5: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4](#), on page 350.

Return to [Configuration Examples for Optimized Layer 2 Overlay Multicast](#), on page 250.

Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric

The following sections provide sample outputs for **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 the 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 12:25:57 UTC

```

Return to [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

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	Gil/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	Gil/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	*	Gil/0/10:101, 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	Gil/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 V1102
```

```
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#!

```

Return to [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

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       Te1/0/10, 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:
  TenGigabitEthernet1/0/10 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    Tel/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, Tel/0/10:101

```

```

Leaf-03# show l2route evpn multicast smet group 239.1.1.1
EVI  ETAG    Origin          Group           Filter         Source(s)
-----
101  0        Tel/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      Te1/0/11, 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:
  TenGigabitEthernet1/0/11 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	Tel/0/11	IGMPv2	N/A	(*, 239.1.1.2)
102	102	Tel/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	*	Tel/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	Tel/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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

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 mapping
PIM Group-to-RP Mappings

Group(s): 224.0.0.0/4, Static
          RP: 192.168.255.255 (?)

EXT# show ip mroute 239.1.1.2
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:

```

Return to [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

Outputs to verify configuration on Spine Switch 1

Verify IGMPv2

```

Spine-01# show bgp 12vpn 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 12vpn 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

Return to [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

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


Return to [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

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 handed off to an external network.

For IPv4 multicast traffic, refer [Figure 30: Topology for Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 Traffic](#).

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

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

enabled. Rendezvous Point (RP) is located outside the EVPN VXLAN fabric. The following tables provide sample configurations for the devices in this topology:

Table 39: Configure VTEP1, CGW, and VTEP3

VTEP1	CGW	VTEP3
-------	-----	-------

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

VTEP1	CGW	VTEP3
<pre> 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 </pre>	<pre> 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 ! </pre>	<pre> 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 TenGigabitEthernet1/0/1 no switchport ip address 172.16.16.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface TenGigabitEthernet1/0/2 no switchport ip address 172.16.26.6 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 </pre>

VTEP1	CGW	VTEP3
<pre> 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 </pre>	<pre> 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 ! </pre>	<pre> ! interface TenGigabitEthernet1/0/10 switchport access vlan 101 switchport mode access ! interface TenGigabitEthernet1/0/11 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 </pre>

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

VTEP1	CGW	VTEP3
	<pre> 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 </pre>	

VTEP1	CGW	VTEP3
	<pre>! end</pre>	

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

Table 40: 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
<pre> Spine-01#show running-config hostname Spine-01 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.13.1 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.14.1 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 exit-address-family </pre>	<pre> Spine-02#show running-config hostname Spine-02 ! ip routing ! 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 GigabitEthernet1/0/1 no switchport ip address 172.16.23.2 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.24.2 255.255.255.0 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 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 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 update-source Loopback0 ! address-family ipv4 exit-address-family </pre>	<pre> EXT1#show running-config hostname EXT ! ip multicast-routing distributed ! ipv6 unicast-routing ipv6 multicast-routing ! 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 ipv6 address 2001:192:168:255::255/128 ipv6 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 ipv6 address 2001:192:168:1::2/64 ipv6 ospf 2 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet0/0/3 no ip address negotiation auto service instance 103 ethernet encapsulation dot1q 103 rewrite ingress tag pop 1 symmetric 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 ipv6 address 2001:192:168:103::1/64 ipv6 ospf 2 area 0 ! interface BDI104 ip address 192.168.104.1 255.255.255.0 </pre>

Spine Switch 1	Spine Switch 2	External Device
<pre> ! 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 </pre>	<pre> ! 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 </pre>	<pre> ip pim sparse-mode ip igmp static-group 239.1.1.2 source 10.1.102.11 ip igmp version 3 ip ospf 2 area 0 ipv6 address 2001:192:168:104::1/64 ipv6 ospf 2 area 0 ! router ospfv3 2 ! address-family ipv6 unicast passive-interface BDI103 passive-interface BDI104 passive-interface Loopback255 exit-address-family ! router ospf 2 passive-interface BDI103 passive-interface BDI104 ! ip pim rp-address 192.168.255.255 ip pim ssm range SSM ! ip access-list standard SSM 10 permit 239.1.1.2 ! ipv6 pim rp-address 2001:192:168:255::255 ! end </pre>

Return to [Example 6: Configuring Optimized Layer 2 Overlay Multicast Interworking with External Network for IPv4 and IPv6](#), on page 377.

Return to [Configuration Examples for Optimized Layer 2 Overlay Multicast](#), on page 250.

Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric

The following sections provide sample outputs for 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 outputs of **show** commands for IPv4 multicast traffic, refer [Verifying Optimized Layer 2 Overlay Multicast for IPv4 Traffic with RP Outside the Fabric](#)

The outputs of **show** commands for IPv6 multicast traffic are the following:

[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 Last-Join/	Interface	Reporter	Vlan	Uptime Last-Leave
::/FF06::239:1:1:1 20:58:46 /	Tu0	FE80::AC10:FE04	101	20:58:46
::/FF06::239:1:1:1 20:48:20 /	Tu0	FE80::AC10:FE06	101	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 12vpn 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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric](#), on page 386.

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 V1101
Y   2001:10:1:101::1              101    101    7c21.0dbd.9541 V1101
```

```
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 Last-Join/	Interface	Reporter	Vlan	Uptime Last-Leave
::FF06::239:1:1:1 00:01:28 /	Gi1/0/10	FE80::46D3:CAFF:FE28:6CC1	101	00:00:00 00:01:28
::FF06::239:1:1:1 20:52:37 /	Tu0	FE80::AC10:FE06	101	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
```

Return to [Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric](#), on page 386.

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:BDFE: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 Last-Join/	Interface	Reporter	Vlan	Uptime	Last-Leave
::/FF06::239:1:1:1 00:00:10 /	Te1/0/10	FE80::EEE1:A9FF:FE37:92C1	101	00:00:00	00:00:10
::/FF06::239:1:1:1 21:05:03 /	Tu0	FE80::AC10:FE04	101	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
Pseudopoports:
  TenGigabitEthernet1/0/10 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    Te1/0/10      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, Te1/0/10:101

```

```

Leaf-03# show l2route evpn multicast smet group FF06::239:1:1:1
EVI  ETAG   Origin          Group          Filter
-----
Source(s)
-----
101  0      Te1/0/10:101    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

```

Return to [Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric](#), on page 386.

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

Return to [Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric](#), on page 386.

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

Return to [Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric](#), on page 386.

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:FEED:9564
  Inherited outgoing interface list:
    BDI103, Forward, 21:07:51/never

```

Return to [Verifying Optimized Layer 2 Overlay Multicast Configuration for IPv4 and IPv6 Traffic with RP Outside the Fabric](#), on page 386.



CHAPTER 10

Configuring Tenant Routed Multicast

- [Restrictions for Tenant Routed Multicast](#) , on page 399
- [Information about Tenant Routed Multicast](#), on page 399
- [How to Configure Tenant Routed Multicast](#), on page 413
- [Verifying Tenant Routed Multicast](#), on page 421
- [Troubleshooting Tenant Routed Multicast](#), on page 422
- [Configuration Examples for Tenant Routed Multicast](#), on page 422

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.

- 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.

356531

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



357605

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



356532

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.

- 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:

- Receiver sends (*,G) IGMP Join to VTEP 3. (*,G) is created at VTEP 3.
- VTEP 3 sends MVPN route type 6 to VTEP 2 which is the RP. (*,G) is created at VTEP 2.
- The source device starts streaming data and (S,G) is created on VTEP 1.
- VTEP 1 performs source-registration at VTEP 2 since it is the RP. (S,G) is created at VTEP 2.
- 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.

356801

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 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



356533

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

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 101](#) for detailed steps to configure Layer 2 and Layer 3 overlay networks.

Perform the following set of procedures to configure TRM in an EVPN VXLAN network:

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 <i>vrf-name</i> Example: Device(config)# vrf definition green	Names the VRF and enters VRF configuration mode.
Step 4	address-family { ipv4 ipv6 } Example: Device(config-vrf)# address-family ipv4	Specifies the VRF and enters VRF address family configuration mode. <ul style="list-style-type: none"> • Use the ipv4 keyword to configure IPv4 address family. • Use the ipv6 keyword to configure IPv6 address family.
Step 5	mdt default vxlan <i>group-address</i> 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.

	Command or Action	Purpose
		<ul style="list-style-type: none"> 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 <i>data-mdt-subnet</i> <i>data-mdt-mask</i> [list <i>access-list-number</i>] 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 <i>kb/s</i> 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.
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 <i>vrf-name</i> Example:	Enables IP multicast forwarding on the overlay VRF.

	Command or Action	Purpose
	Device(config)# ip multicast-routing vrf green	
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.

	Command or Action	Purpose
Step 6	interface vlan <i>access-facing-vlan-id</i> 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 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 <i>autonomous-system-number</i> 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. <ul style="list-style-type: none"> • 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.

	Command or Action	Purpose
Step 6	neighbor <i>ip-address</i> 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 <i>ip-address</i> advertisement-interval <i>seconds</i> 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 <i>ip-address-of-rp</i> Example: Device(config)# ip pim rp-address <rp-ip-address>	Configures the RP in the underlay network. For information about RP redundancy, see <i>IP Multicast Routing Configuration Guide</i> .

	Command or Action	Purpose
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 <i>loopback-interface</i> Example: Device(config)# interface Loopback 13	Enters interface configuration mode for the specified Loopback interface.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# vrf forwarding green	Configures forwarding table for the Loopback interface.
Step 5	ip-address <i>ip-address subnet-mask</i> 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.

	Command or Action	Purpose
Step 8	<pre>{ip ipv6 } pim vrf vrf-name rp-address rp-address</pre> <p>Example:</p> <pre>Device(config)# ip pim vrf green rp-address 10.1.13.13</pre>	<p>Configures the address of the local VTEP as the PIM RP for the multicast group.</p> <ul style="list-style-type: none"> • 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. <p>Note The loopback interface specified must be part of the same VRF.</p>
Step 9	<pre>{ip ipv6 } pim vrf vrf-name register-source loopback-address-of-vtep</pre> <p>Example:</p> <pre>Device(config)# ip pim vrf green register-source loopback901</pre>	<p>Configures a unique IP address for the loopback interface of the VTEP that acts as the first hop router to multicast traffic.</p>
Step 10	<pre>end</pre> <p>Example:</p> <pre>Device(config)# end</pre>	<p>Returns to privileged EXEC mode.</p>

Configuring TRM with PIM Source Specific Mode

To configure TRM with PIM Source Specific Mode, perform the following tasks:

- [Configuring the TRM Multicast Distribution Tree in the VRF](#)
- [Configuring Multicast Routing on the Overlay VRF](#)
- [Configuring Multicast on Switch Virtual Interfaces for Core-facing and Access-facing VLANs](#)
- [Configuring 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	<pre>enable</pre> <p>Example:</p> <pre>Device> enable</pre>	<p>Enables privileged EXEC mode.</p> <p>Enter your password, if prompted.</p>

	Command or Action	Purpose
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.
Step 3	ip pim [vrf vrf-name] ssm {default range access-list } Example: Device(config)# <code>ip pim vrf green ssm default</code>	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)# <code>end</code>	Returns to privileged EXEC mode.

Verifying Tenant Routed Multicast

The following table lists the **show** commands that are used to verify TRM:

Command	Purpose
<code>show nve peers</code>	Displays NVE interface state information for peer leaf switches.
<code>show l2vpn evpn peers vxlan</code>	Displays Layer 2 EVPN peer route counts in the VXLAN and up time.
<code>show ip igmp vrf green groups</code>	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.
<code>show bgp ipv4 mvpn all</code>	Displays the MVPN options for BGP MVPN C-route signaling.
<code>show ip mroute vrf green</code>	Displays the contents of the mroute table that pertain to a specific MVRF instance.
<code>show ip mfib vrf green</code>	Displays forwarding entries and interfaces in the IPv4 Multicast Forwarding Information Base (MFIB) associated with MVRF instances.
<code>show ip mroute</code>	Displays multicast routing table information.
<code>show ip mfib</code>	Displays the forwarding entries and interfaces in the IPv4 MFIB.

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.

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 41: 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
<pre>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 !</pre>	<pre>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 !</pre>	<pre>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 !</pre>

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

VTEP 1	VTEP 2	VTEP 3
<pre> 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 ! </pre>	<pre> 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 ! </pre>	<pre> 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 TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/10 switchport access vlan 102 switchport mode access ! </pre>

VTEP 1	VTEP 2	VTEP 3
<pre> 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 ! </pre>	<pre> 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 ! </pre>	<pre> 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 172.16.255.6 ! 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 ! </pre>

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

VTEP 1	VTEP 2	VTEP 3
<pre> 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# </pre>	<pre> 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# </pre>	<pre> 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# </pre>

Table 42: 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
<pre> 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 ! </pre>	<pre> 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 ! </pre>

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

Spine Switch 1	Spine Switch 2
<pre> 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# </pre>	<pre> 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# </pre>

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 429](#)
- [Outputs to Verify the Configuration on VTEP 2, on page 436](#)
- [Outputs to Verify the Configuration on VTEP 3, on page 443](#)
- [Outputs to Verify the Configuration on Spine Switch 1, on page 451](#)

- [Outputs to Verify the Configuration on Spine Switch 2, on page 455](#)

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
```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

::: 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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

          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
```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

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 - encaps-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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```
(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 428](#)

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 rinfo 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 rinfo 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 ?
```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

*>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 ?
*> [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 ?
*> [5] [1:1] [0] [128] [FC00:2:255::255]/29
::                      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:

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```
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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

(*, 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 428](#)

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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

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 [2] [172.16.254.3:101] [0] [48] [10B3D56A8FC1] [128] [FC00:1:101::1]/36
      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 [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [32] [10.1.101.11]/24
      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 [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
      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 ?
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 [2] [172.16.254.4:102] [0] [48] [44D3CA286CC5] [32] [10.1.102.12]/24
      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 [2] [172.16.254.4:102] [0] [48] [44D3CA286CC5] [128] [FE80::46D3:CAFF:FE28:6CC5]/36
      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 [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [128] [FC00:1:102::1]/36

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

          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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: SJCx
  Incoming interface: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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 428

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 rinfo 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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

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 ?

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

* 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
Tunnell 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
Tunnell 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
Tunnell 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 428](#)

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
```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

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 ?

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

*>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

```

Example: Configuring TRM in PIM Sparse Mode with Anycast RP

```

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
Tunnell 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
Tunnell 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
Tunnell 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
Tunnell 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 428

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 43: 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
<pre>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</pre>	<pre>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 !</pre>	<pre>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 !</pre>

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

VTEP 1	Border VTEP	VTEP 2
<pre> ! 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 ! 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 ! 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 ! </pre>	<pre> 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 TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/5 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 ! </pre>	<pre> 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 GigabitEthernet1/0/10 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 </pre>

VTEP 1	Border VTEP	VTEP 2
<pre> 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 </pre>	<pre> 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 </pre>	<pre> ! 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 ! </pre>

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

VTEP 1	Border VTEP	VTEP 2
<pre> ! 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# </pre>	<pre> ! 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# </pre>	<pre> 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# </pre>

Table 44: 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
<pre> 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 ! </pre>	<pre> 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 ! </pre>

Spine Switch 1	Spine Switch 2
<pre> 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# </pre>	<pre> 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# </pre>

Table 45: 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
<pre> 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 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# </pre>	<pre> 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# </pre>	<pre> 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# </pre>

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 \(RP Inside the BGP EVPN VXLAN Fabric\)](#)
- [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
```

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

```

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: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse, 1d05h/00:02:10
```

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

```

(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 469](#)

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
```

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

```
(* , 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 - encaps-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
```

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

```
(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 469

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_RT's  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
```

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

```

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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, Forward/Sparse, 1d05h/00:02:43, A

(*, 225.0.0.101), 1d05h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, RPF nbr 172.16.26.2

```

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

```

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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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

```

```

TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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 469

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

```

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

```

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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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
Tunnell 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
Tunnell 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
Tunnell 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
  Tunnell 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
  Tunnell 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
  Tunnell 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 469

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
```

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

```

(*, 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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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: GigabitEthernet1/0/2, 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: GigabitEthernet1/0/4, 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
  Tunnell 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
  Tunnell 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
  Tunnell 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
  Tunnell 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 469

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 [Figure 43: 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 46: 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
<pre>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 !</pre>	<pre>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 !</pre>	<pre>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 !</pre>

Example: Configuring 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
<pre> 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 </pre>	<pre> 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 TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/5 switchport trunk allowed vlan 2001 switchport mode trunk ! interface TenGigabitEthernet1/0/10 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 ! </pre>	<pre> 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 </pre>

VTEP 1	Border VTEP	VTEP 2
<pre> ! 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 ! 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 ! </pre>	<pre> 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 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 </pre>	<pre> ! 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 ! 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 </pre>

Example: Configuring 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
<pre> 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# </pre>	<pre> ! 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 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# </pre>	<pre> ! 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# </pre>
	<pre> ! 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# </pre>	

Table 47: 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 Switch 2
<pre> 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 ! </pre>	<pre> 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 ! </pre>

Spine Switch 1	Spine Switch 2
<pre> 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# </pre>	<pre> 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# </pre>

Table 48: 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
<pre> 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 ! </pre>	<pre> 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 </pre>	<pre> 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 ! </pre>

Router 1	Router 2	Router 3
<pre>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#</pre>	<pre>! ipv6 pim rp-address FC00:2:255::255 ! end ! R2#</pre>	<pre>ipv6 pim rp-address FC00:2:255::255 ! end ! R3#</pre>

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 \(RP Inside the BGP EVPN VXLAN Fabric\)](#)
- [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)

FF06:1::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 502

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: V1901,::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 502

Outputs to Verify the Configuration on Border VTEP

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

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

```

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 rinfo 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: V1901,::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] [1:1] [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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
  TenGigabitEthernet1/0/2, Forward/Sparse, 1d06h/00:03:21, A

(*, 225.0.0.101), 1d06h/stopped, RP 172.16.255.255, flags: SJCx
Incoming interface: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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

```

```

TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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 502

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 rinfo 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

```

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

```

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 502

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 rinfo 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,

```

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

```

    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
Tunnell 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
Tunnell 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
Tunnell 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
Tunnell 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 502

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:

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

Table 49: 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
<pre> 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 ! </pre>	<pre> 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 ! </pre>	<pre> 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 ! </pre>

VTEP 1	Border VTEP	VTEP 2
<pre> 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 ! </pre>	<pre> interface TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/5 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 </pre>	<pre> 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 ! </pre>

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

VTEP 1	Border VTEP	VTEP 2
<pre> 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# </pre>	<pre> ! 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# </pre>	<pre> 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# </pre>

Table 50: 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 Switch 2
<pre> 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 ! </pre>	<pre> 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 ! </pre>

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

Spine Switch 1	Spine Switch 2
<pre> 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# </pre>	<pre> 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# </pre>

Table 51: 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
<pre> 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# </pre>	<pre> 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# </pre>	<pre> 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# </pre>

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 540](#)
- [Outputs to Verify the Configuration on VTEP 2, on page 545](#)
- [Outputs to Verify the Configuration on Border VTEP, on page 550](#)
- [Outputs to Verify the Configuration on Spine Switch 1, on page 556](#)
- [Outputs to Verify the Configuration on Spine Switch 2, on page 559](#)
- [Outputs to Verify the Configuration on Router 3 \(RP Outside the BGP EVPN VXLAN Fabric\), on page 563](#)

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
```

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

```
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

```

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

```

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 539

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: SJPLGx
  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

```

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

```
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
```

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

```

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 539](#)

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 rinfo 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/stoppped, 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/stoppped

(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/stoppped

(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   [5][1:1][10.1.102.12][226.1.1.1]/18
      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   [6][1:1][65001][10.2.255.255/32][224.0.1.40/32]/22
      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   [6][1:1][65001][10.2.255.255/32][226.1.1.1/32]/22
      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 ?

```

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

```

* 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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 539](#)

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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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: 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
  Tunnell 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
  Tunnell 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
  Tunnell 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
  Tunnell 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
  Tunnell 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 539

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
```

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

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:
  GigabitEthernet1/0/4, Forward/Sparse, 00:44:04/00:02:52
  GigabitEthernet1/0/1, 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
```

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

```

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
  Tunnell 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
  Tunnell 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
  Tunnell 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
Tunnell 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 539

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:
    TenGigabitEthernet0/0/0, Forward/Sparse, 17:10:20/00:03:17
    TenGigabitEthernet0/0/1, Forward/Sparse, 2d16h/00:03:11

(10.2.255.1, 226.1.1.1), 00:37:40/00:02:14, flags: PJT
  Incoming interface: TenGigabitEthernet0/0/0, 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: TenGigabitEthernet0/0/0, 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
    TenGigabitEthernet0/0/1, Forward/Sparse, 2d19h/00:03:26
    TenGigabitEthernet0/0/0, 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 539](#)

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 [Figure 44: 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 52: 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
<pre> 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 </pre>	<pre> 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 ! </pre>	<pre> 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 ! </pre>

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

VTEP 1	Border VTEP	VTEP 2
<pre> ! 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 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 ! 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 ! </pre>	<pre> 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 TenGigabitEthernet1/0/1 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/5 switchport trunk allowed vlan 2001 switchport mode trunk ! interface TenGigabitEthernet1/0/10 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 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 ! </pre>	<pre> 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 ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable no autostate ! </pre>

VTEP 1	Border VTEP	VTEP 2
<pre> 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 ! </pre>	<pre> 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::1/64 ipv6 enable 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 </pre>	<pre> 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 ! </pre>

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

VTEP 1	Border VTEP	VTEP 2
<pre> 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# </pre>	<pre> ! 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# </pre>	<pre> 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# </pre>

Table 53: 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
----------------	----------------

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

Spine Switch 1	Spine Switch 2
<pre> 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 ! </pre>	<pre> 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 ! </pre>

Spine Switch 1	Spine Switch 2
<pre> 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# </pre>	<pre> 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# </pre>

Table 54: 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
<pre> 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 </pre>	<pre> 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 </pre>	<pre> 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 ! </pre>

Router 1	Router 2	Router 3
<pre> ! 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# </pre>	<pre> ! ip pim rp-address 10.2.255.255 ! ipv6 pim rp-address FC00:2:255::255 ! end ! R2# </pre>	<pre> 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# </pre>

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 573](#)
- [Outputs to Verify the Configuration on VTEP 2, on page 581](#)
- [Outputs to Verify the Configuration on Border VTEP, on page 588](#)
- [Outputs to Verify the Configuration on Spine Switch 1, on page 596](#)
- [Outputs to Verify the Configuration on Spine Switch 2, on page 599](#)
- [Outputs to Verify the Configuration on Router 3 \(RP Outside the BGP EVPN VXLAN Fabric\), on page 603](#)

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 rinfo 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
IP PIM Group Mapping Table
(* indicates group mappings being used)

FF00::/8*
  SM, RP: FC00:2:255::255
  RPF: V1901,::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
```

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

```

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

```

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

```

(*,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-stargy 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,

```

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

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 573](#)

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)

FF06:1::/8*
  SM, RP: FC00:2:255::255
  RPF: V1901,::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
      ::
      * i [5] [1:1] [FC00:2:255::1] [FF06:1::1] /42

```

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

```

          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

```

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

```

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 573

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 rinfo 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: V12001, 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   [5][1:1][FC00:2:255::1][FF06:1::1]/42
      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   [6][1:1][65001][FC00:2:255::1][FF06:1::1]/46
      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   [7][1:1][65001][FC00:2:255::1][FF06:1::1]/46
      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,

```

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

```

    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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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:
    TenGigabitEthernet1/0/2, Forward/Sparse, 1d01h/00:02:49, A

(*, 225.0.0.101), 1d01h/stopped, RP 172.16.255.255, flags: SJCx
  Incoming interface: TenGigabitEthernet1/0/2, 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: TenGigabitEthernet1/0/2, 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
  TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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
TenGigabitEthernet1/0/2 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 573

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 [5] [1:1] [FC00:2:255::1] [FF06:1::1] /42
      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: GigabitEthernet1/0/4, 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: GigabitEthernet1/0/2, 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: 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
  Tunnell 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
  Tunnell 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
  Tunnell 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
  Tunnell 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
Tunnell 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 573

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 rinfo 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 ?

```

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

```

*>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,

```

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

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
Tunnell 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
Tunnell 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
Tunnell 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
Tunnell 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 573](#)

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:
    TenGigabitEthernet0/0/1, Forward, 1d04h/00:03:12
    TenGigabitEthernet0/0/0, 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:
    TenGigabitEthernet0/0/0, Null, 00:18:43/00:02:45
    TenGigabitEthernet0/0/1, 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:
    TenGigabitEthernet0/0/1, 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:
    TenGigabitEthernet0/0/0, Null, 00:18:08/00:02:45
    TenGigabitEthernet0/0/1, Null, 00:18:04/00:03:12

(FC00:2:255::1, FF06:1::1), 00:18:06/00:02:55, flags: S
  Incoming interface: TenGigabitEthernet0/0/0
  RPF nbr: FE80::A2B4:39FF:FE21:9181
  Inherited Outgoing interface list:
    TenGigabitEthernet0/0/1, 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 573



CHAPTER 11

Configuring EVPN VXLAN External Connectivity

- [Restrictions for EVPN VXLAN External Connectivity, on page 605](#)
- [Information About EVPN VXLAN External Connectivity, on page 605](#)
- [How to Configure EVPN VXLAN External Connectivity, on page 610](#)
- [Configuration Examples for EVPN VXLAN External Connectivity, on page 624](#)

Restrictions for EVPN VXLAN External Connectivity

- External connectivity with VPLS networks is supported only when bridging is the mode of interworking between the two domains. Integrated routing and bridging (IRB) is not supported between a BGP EVPN VXLAN fabric and a VPLS network.
- External Connectivity with Layer 3 networks is supported only for IPv4 and IPv6 unicast traffic.
- External connectivity with an MVPN network is not supported for multicast traffic.
- Import of EVPN IP routes, which includes both route type 5 and route type 2 host routes, to global routing table is not supported.
- not support the following features:
 - MPLS Layer 3 VPN Border Leaf Handoff
 - MPLS Layer 3 VPN Border Spine Handoff
 - VPLS over MPLS Border Leaf Handoff
 - VPLS over MPLS Border Spine Handoff
 - Interworking of Layer 3 TRM with MVPN Networks for IPv4 Traffic

Information About EVPN VXLAN External Connectivity

External connectivity allows the movement of Layer 2 and Layer 3 traffic between an EVPN VXLAN network and an external network. It also enables the EVPN VXLAN network to exchange routes with the externally connected network. Routes within an EVPN VXLAN network are already shared between all the VTEPs or leaf switches. External connectivity uses the VTEPs on the periphery of the network to pass on these routes to an external Layer 2 or Layer 3 network. Similarly, the EVPN VXLAN network imports the reachability

routes from the external network. External connectivity extends the Layer 2 or Layer 3 overlay network outside the VXLAN network. The process of extending a Layer 2 or Layer 3 network outside the EVPN VXLAN network is also known as handoff.

Implementation of Border Nodes for EVPN VXLAN External Connectivity

Border nodes or border VTEPs are the devices through which you establish a connection between an EVPN VXLAN network and an external network. The border nodes sit on the periphery of the EVPN VXLAN network and remain a part of the BGP EVPN VXLAN fabric. To enable external connectivity, you can implement the border nodes of an EVPN VXLAN network as either border leaf or border spine switches.

Connectivity Through a Border Leaf Switch


Leaf switches deployed as border nodes support the required control plane and data plane functionalities. Border leaf deployment ensures that the configuration on the spine switches is much simpler. Border leaf switches only allow communication between the external network and the VXLAN network, also known as north-south communication.



Note A border leaf switch can also be multiple switches functioning as a single logical system with Cisco StackWise Virtual configured.

The following figure shows border leaf external connectivity of an EVPN VXLAN network with external Layer 2 and Layer 3 networks:

Figure 45: EVPN VXLAN External Connectivity Through a Border Leaf Switch




Connectivity Through a Border Spine Switch

Deploying spine switches as border nodes provides the advantage of optimizing the north-south communication with external resources. At the same time, border spine deployment allows the spine switches to support VXLAN control and data plane functionality. Border spine switches allow both north-south communication and east-west communication. East-west communication represents the communication within the nodes of the EVPN VXLAN network.

The following figure shows border spine external connectivity of an EVPN VXLAN network with external Layer 2 and Layer 3 networks.:

Figure 46: EVPN VXLAN External Connectivity Through a Border Spine Switch



356587

External Connectivity with Layer 3 Networks

Layer 3 external connectivity or handoff is established by connecting the border nodes of a BGP EVPN VXLAN fabric with an edge router from the external Layer 3 network. The border node acts as a VTEP to perform VXLAN encapsulation and decapsulation, but it also routes the traffic towards the edge routing device. The VXLAN-facing interface on the external Layer 3 network can be a switch virtual interface (SVI), or a Layer 3 interface, or a Layer 3 subinterface.

You can use Layer 3 external connectivity to achieve any of the following:

- Extend the logical isolation between VRFs or VLANs within the EVPN VXLAN network into the externally routed network. The external routed network can be a traditional non-VXLAN campus network, a datacenter, or a WAN.
- Provide shared access within the EVPN VXLAN network to a common external service such as the internet.

BGP EVPN VXLAN fabric supports Layer 3 external connectivity with VRF-Lite and MPLS Layer 3 VPN networks.

Layer 3 External Connectivity with VRF-Lite

Using VRF allows for the use of multiple routing tables that are independent and isolated. VRF-Lite is a mechanism to extend the tenant Layer 3 VRF information beyond the BGP EVPN VXLAN Fabric. External connectivity with VRF-Lite or VRF handoff involves a two-box approach where the border node and the edge router are physically independent devices. With VRF-Lite handoff, the BGP EVPN VXLAN fabric extends the connectivity for different tenants externally on a hop-by-hop basis.

Once the border node learns external routes from the edge router, it advertises the prefixes inside the BGP EVPN VXLAN fabric as EVPN type 5 routes. This information is distributed to all the other VTEPs in the network. The border node also advertises EVPN routes to the external edge router. It sends the EVPN routes learned from the Layer 2 VPN EVPN address family to the IPv4 or IPv6 unicast address family.

Layer 3 Multicast External Connectivity with MPLS Layer 3 VPN

Layer 3 external connectivity with an MPLS Layer 3 VPN network or MPLS handoff uses a single-box approach. The single-box approach combines the functionalities of an EVPN VXLAN border node and an MPLS PE router into a single physical device. The device is also known as a border PE node. The border PE node reoriginates IP prefixes from the EVPN address family of the BGP EVPN VXLAN fabric to the VPNv4 address family of the MPLS network. Likewise, the border PE node performs the corresponding function in the reverse direction. eBGP peering is necessary between the border PE node and the MPLS PE devices to ensure the connectivity.

MPLS handoff allows scalability for EVPN VXLAN networks that have a large number of tenants or VRFs. Scalability is not possible with VRF-Lite handoff.

In every VRF on a border VTEP, there are two sets of manually configured import and export route targets. The first set of import and export route targets is associated with the BGP neighbor in the BGP EVPN VXLAN fabric. This BGP neighbor uses the EVPN address family to exchange Layer 3 information. The second set of import and export route targets is associated with the BGP neighbor in the Layer 3 VPN network. This BGP neighbor uses either VPNv4 or VPNv6 unicast address families to exchange Layer 3 information. The separation of route targets allows you to configure both sets of route targets independently. In this way, a border VTEP in an EVPN VXLAN network effectively stitches the two sets of route targets. The route targets associated with the BGP neighbor in the Layer 3 VPN network are known as normal route targets. The route targets associated with the BGP neighbor in the BGP EVPN VXLAN fabric are known as stitching route targets.

External Connectivity with Layer 2 Networks

Layer 2 external connectivity or handoff for an EVPN VXLAN network extends the Layer 2 domain outside of the network. BGP EVPN VXLAN fabric supports Layer 2 external connectivity with IEEE 802.1Q, access, and VPLS over MPLS networks.

Layer 2 External connectivity with IEEE 802.1Q or Access Networks

Layer 2 handoff to IEEE 802.1Q networks is achieved through a regular IEEE 802.1Q Trunk port configuration on the Switchport interfaces on the border nodes. You can also connect EVPN VXLAN networks to external access networks.

The commonly deployed scenario has EVPN enabled at the distribution layer and has the access layer switches connected with IEEE 802.1Q Trunk encapsulation. The IEEE 802.1Q Layer 2 traffic that comes from the access layer switches is mapped to the corresponding VLAN. The border node then bridges the traffic towards the destination with VXLAN encapsulation. The inner packet does not carry the IEEE 802.1Q tag. Instead, the VXLAN network identifier (VNI), which is the Layer 2 VNI in the VXLAN header, represents the broadcast

domain. Similarly, the border nodes decapsulate the traffic from the BGP EVPN VXLAN fabric and bridge it with the corresponding IEEE 802.1Q tag to the access switches. The interface on the border VTEP that faces the external interface can be either an access or a Trunk port. The external interface can belong to either a Layer 2 switch or a firewall.



Note If you connect the network to an external Layer 2 switch through two border VTEPs, it represents a dual connection. In such cases, STP does not propagate over the BGP EVPN VXLAN fabric by default.

Layer 2 External connectivity with VPLS over MPLS Network

External connectivity with VPLS networks or VPLS handoff is achieved when a border VTEP or multiple border VTEPs establish a connection with the VPLS network. The border nodes act as the provider edge (PE) devices in the VPLS network and as VTEPs in the EVPN VXLAN network.

BGP EVPN VXLAN supports VPLS handoff in the form of VPLS stitching through either an access VFI or an access pseudowire on the VLAN on the border VTEP.

The access pseudowires and the pseudowires in the access VFI function as the access ports in the EVPN VXLAN network. The BGP EVPN VXLAN fabric treats the MAC addresses learned on the pseudowires as locally learned MAC addresses. It advertises these MAC addresses within the fabric as EVPN type 2 routes. The pseudowires are in a different split horizon group compared to the EVPN VXLAN network. Therefore, BUM traffic floods between both the EVPN VXLAN and VPLS networks.

Route Leaking between EVPN-enabled VRF and Global Routing Tables

The route leaking feature provides the import and export of routes between an EVPN-enabled Virtual Routing and Forwarding (VRF) Table and the Global Routing Table (GRT). By enabling the import or export of the EVPN learned routes into the GRT, the border node provides access to services such as DNS or DHCP to the hosts in the fabric. Routes that are imported from GRT to VRF can subsequently be imported to EVPN using the **import {ipv4|ipv6} unicast map allow-evpn** command. Routes that are imported from EVPN to IP VRF can subsequently be exported to GRT using the **export {ipv4|ipv6} unicast map allow-evpn** command.

You can further filter the imported routes using route-maps. A route-map contains the match conditions that define which of the routes from the import or export tables are allowed into the target routing table. This reduces the risk of unintended import or export of routes between the tables and helps contain the table sizes and advertisements.

How to Configure EVPN VXLAN External Connectivity

This section provides information about how to configure external connectivity between an EVPN VXLAN network and an external Layer 2 or Layer 3 network.

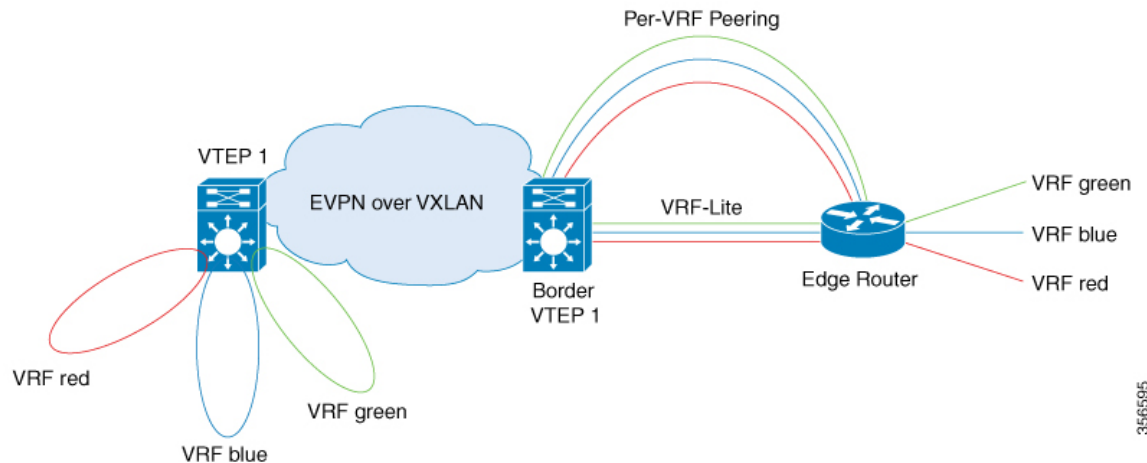


Note You must configure EVPN VXLAN Layer 2 and Layer 3 overlay networks before you configure external connectivity. See [How to Configure EVPN VXLAN Integrated Routing and Bridging](#) for detailed steps.

Enabling Layer 3 External Connectivity with VRF-Lite

The following figure shows a sample topology that illustrates Layer 3 external connectivity with VRF-Lite:

Figure 47: Layer 3 External Connectivity with VRF-Lite



356595

To configure Layer 3 external connectivity with VRF-Lite, perform the following set of procedures:

- Configure the VRF on the border VTEP interface that faces the external router.
- Ensure that Layer 2 VPN EVPN is advertised as part of the BGP VRF configuration. See [Configuring BGP with EVPN and VRF Address Families on a VTEP, on page 105](#) for detailed steps.



Note Redistribution of the respective interior gateway protocol (IGP) is required in the BGP VRF address family to distribute the external prefixes into the BGP EVPN VXLAN fabric.

For more information about VRF-Lite, see *Contents* → *IP Routing Configuration Guide* → *Configuring VRF-lite* in the software configuration guide for the applicable release.

Configuring the VRF on the Border VTEP Interface that Faces the External Router

To configure the VRF on the border VTEP interface that faces the external router, perform these 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# <code>configure terminal</code>	Enters global configuration mode.
Step 3	interface <i>interface-id</i> Example: Device(config)# <code>interface GigabitEthernet1/0/30</code>	Enters the interface configuration mode for the specified interface.
Step 4	vrf forwarding <i>vrf-name</i> Example: Device(config-if)# <code>vrf forwarding green</code>	Associates the VRF with the interface. Note 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 address <i>ip-address</i> Example: Device(config-if)# <code>ip address 192.168.3.203 255.255.255.0</code>	Configures the IP address for the interface.
Step 6	end Example: Device(config-if)# <code>end</code>	Returns to privileged EXEC mode.

Enabling Route Leaking between EVPN-enabled VRF and Global Routing Table

Route Leaking between an EVPN-enabled VRF and the Global Routing Table (GRT) is required to access shared services such as DNS or DHCP, which are reachable through the global routing table.

To configure route leaking on a leaf node, perform the following steps.

Procedure

	Command or Action	Purpose
Step 1	enable Example: Device> <code>enable</code>	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	configure terminal Example: Device# <code>configure terminal</code>	Enters global configuration mode.

	Command or Action	Purpose
Step 3	vrf definition <i>vrf-name</i> Example: Device(config)# vrf definition green	Creates a VRF instance.
Step 4	address-family { ipv4 ipv6 } Example: Device(config-vrf)# address-family ipv4	Configures the IPv4 or IPv6 address family and enters the VRF address family configuration mode. <ul style="list-style-type: none"> • Use the ipv4 keyword to configure IPv4 address family. • Use the ipv6 keyword to configure IPv6 address family.
Step 5	import { ipv4 ipv6 } unicast map <i>import-route-map allow-evpn</i> Example: Device(config-vrf-af)# import ipv4 unicast map RM-GL-TO-EVPN allow-evpn	Specifies an import map to import IPv4 or IPv6 prefixes from the global routing table to a VRF table. Specify ipv4 or ipv6 , which you specified in Step 4. The map import-route-map keyword specifies the route-map that defines the prefixes to import. In the example, RM-GL-TO-EVPN is the name of the route-map that defines the conditions and the IPv4 prefixes to import.
Step 6	export { ipv4 ipv6 } unicast map <i>export-route-map allow-evpn</i> Example: Device(config-vrf-af)# export ipv4 unicast map RM-EVPN-TO-GL allow-evpn	Specifies a route map to export IPv4 or IPv6 prefixes to the global routing table. Specify ipv4 or ipv6 , which you specified in Step 4. The map export-route-map keyword specifies the route-map that defines the prefixes to export. In the example, RM-EVPN-TO-GL is the name of the route-map that defines the conditions and the IPv4 prefixes to export.
Step 7	exit-address-family Example: Device(config-vrf-af)# exit-address-family	Exits VRF address family configuration mode and returns to VRF configuration mode.

Example: Enable Route Leaking Between EVPN and Global Routing Table

The following snippet shows a sample configuration for route leaking on a Leaf VTEP.

```


Leaf-01#
!
ip prefix-list EVPN-TO-GL seq 5 permit 10.1.101.0/24 le 32
ip prefix-list EVPN-TO-GL seq 10 permit 10.1.102.0/24 le 32
!
ip prefix-list GL-TO-EVPN seq 5 permit 192.168.255.0/24 le 32
!
route-map RM-GL-TO-EVPN permit 10
 match ip address prefix-list GL-TO-EVPN
!
route-map RM-EVPN-TO-GL permit 10
 match ip address prefix-list EVPN-TO-GL
!
vrf definition green
 description green VRF definition
 rd 1:1
!
 address-family ipv4
  import ipv4 unicast map RM-GL-TO-EVPN allow-evpn
  export ipv4 unicast map RM-EVPN-TO-GL allow-evpn
  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
!
!
interface Loopback100
 ip address 192.168.255.1 255.255.255.255
 ip ospf 1 area 0
!
interface Vlan101
 vrf forwarding green
 ip address 10.1.101.1 255.255.255.0
!
interface Vlan102
 vrf forwarding green
 ip address 10.1.102.1 255.255.255.0
!

```

Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN

The following figure shows a sample topology that illustrates Layer 3 external connectivity with an MPLS Layer 3 VPN network:

Figure 48: Layer 3 External Connectivity with MPLS Layer 3 VPN



To enable EVPN VLAN Layer 3 external connectivity with MPLS Layer 3 VPN networks, perform the following set of procedures:

- Run the **mpls label mode all-vrfs protocol all-afs per-vrf** command in global configuration mode on the border VTEP.
- Configure BGP with reorigination of routes with a new route type for Layer 2 VPN, VPNv4, VPNv6 address families on the border VTEP.

Configuring BGP on a Border VTEP for External Connectivity with MPLS Layer 3 VPN

To configure BGP on a border VTEP to establish external connectivity with an MPLS Layer 3 VPN network, perform this procedure:

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	router bgp <i>autonomous-system-number</i> 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 section of the <i>IP Routing Configuration Guide</i> .
Step 5	bgp update-delay <i>time-period</i> 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 section of the <i>IP Routing Configuration Guide</i> .
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 section of the <i>IP Routing Configuration Guide</i> .
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 section of the <i>IP Routing Configuration Guide</i> .
Step 8	neighbor <i>spine-ip-address</i> remote-as <i>number</i> Example: Device (config-router) # neighbor 172.16.255.1 remote-as 1	Defines multiprotocol-BGP neighbors in the EVPN network. Use the IP address of the spine switch as the neighbor IP address. This configures the spine switch as a BGP neighbor.
Step 9	neighbor <i>mpls-peer-ip-address</i> remote-as <i>number</i> Example: Device (config-router) # neighbor 172.16.255.103 remote-as 1	Defines multiprotocol-BGP neighbors in the external MPLS network. Use the IP address of the external MPLS network peer as the neighbor IP address. This configures the external MPLS network peer as a BGP neighbor.
Step 10	neighbor { <i>ip-address</i> <i>group-name</i> } update-source <i>interface</i> Example: Device (config-router) # neighbor 172.16.255.1 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.

	Command or Action	Purpose
Step 11	address-family l2vpn evpn Example: Device(config-router)# address-family l2vpn evpn	Specifies the L2VPN address family and enters address family configuration mode.
Step 12	import vpnv4 unicast re-originate Example: Device(config-router-af)# import vpnv4 unicast re-originate	Reoriginates the VPNv4 routes imported from the external peer into the EVPN address family as EVPN routes, and distributes within the EVPN fabric.
Step 13	import vpnv6 unicast re-originate Example: Device(config-router-af)# import vpnv6 unicast re-originate	Reoriginates the VPNv6 routes imported from the external peer into the EVPN address family as EVPN routes, and distributes within the EVPN fabric.
Step 14	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 15	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. Note Use either of extended or both keywords. External connectivity cannot be established when you use the standard keyword.
Step 16	neighbor {ip-address peer-group-name} next-hop-self [all] Example: Device(config-router-af)# neighbor ip-address next-hop-self all	Configures the router as the next hop for a BGP-speaking neighbor or peer group. The all keyword is mandatory when implementing external connectivity through iBGP, where the EVPN fabric and the MPLS network are in the same BGP autonomous system number. The all keyword is optional when implementing external connectivity through eBGP, where the EVPN fabric and the MPLS network are in different BGP autonomous system numbers.
Step 17	exit-address-family Example: Device(config-router-af)# exit-address-family	Exits address family configuration mode and returns to router configuration mode.

	Command or Action	Purpose
Step 18	address-family vpv4 Example: Device(config-router)# address-family vpv4	Specifies the VPNv4 address family and enters address family configuration mode.
Step 19	import l2vpn evpn re-originate Example: Device(config-router-af)# import l2vpn evpn re-originate	Reoriginates the EVPN routes imported from the EVPN fabric into the VPNv4 address family as VPNv4 routes and distributes them to the external network.
Step 20	neighbor ip-address activate Example: Device(config-router-af)# neighbor 172.16.255.103 activate	Enables the exchange information from a BGP neighbor. Use the IP address of the external MPLS network router as the neighbor IP address.
Step 21	neighbor ip-address send-community [both extended standard] Example: Device(config-router-af)# neighbor 172.16.255.103 send-community both	Specifies the communities attribute sent to a BGP neighbor. Use the IP address of the external MPLS network router as the neighbor IP address. Note Use either of extended or both keywords. External connectivity cannot be established when you use the standard keyword.
Step 22	neighbor {ip-address peer-group-name} next-hop-self [all] Example: Device(config-router-af)# neighbor ip-address next-hop-self all	Configures the router as the next hop for a BGP-speaking neighbor or peer group. The all keyword is mandatory when implementing external connectivity through iBGP, where the EVPN fabric and the MPLS network are in the same BGP autonomous system number. The all keyword is optional when implementing external connectivity through eBGP, where the EVPN fabric and the MPLS network are in different BGP autonomous system numbers.
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	address-family vpv6 Example: Device(config-router)# address-family vpv6	Specifies the VPNv6 address family and enters address family configuration mode.

	Command or Action	Purpose
Step 25	import l2vpn evpn re-originate Example: Device(config-router-af)# import l2vpn evpn re-originate	Reoriginates the EVPN routes imported from the EVPN fabric into the VPNv6 address family as VPNv6 routes and distributes them to the external network.
Step 26	neighbor ip-address activate Example: Device(config-router-af)# neighbor 172.16.255.103 activate	Enables the exchange information from a BGP neighbor. Use the IP address of the spine switch as the neighbor IP address.
Step 27	neighbor ip-address send-community [both extended standard] Example: Device(config-router-af)# neighbor 172.16.255.103 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. Note Use either of extended or both keywords. External connectivity cannot be established when you use the standard keyword.
Step 28	neighbor {ip-address peer-group-name} next-hop-self [all] Example: Device(config-router-af)# neighbor ip-address next-hop-self all	Configures the router as the next hop for a BGP-speaking neighbor or peer group. The all keyword is mandatory when implementing external connectivity through iBGP, where the EVPN fabric and the MPLS network are in the same BGP autonomous system number. The all keyword is optional when implementing external connectivity through eBGP, where the EVPN fabric and the MPLS network are in different BGP autonomous system numbers.
Step 29	exit-address-family Example: Device(config-router-af)# exit-address-family	Exits address family configuration mode and returns to router configuration mode.
Step 30	end Example: Device(config-router)# end	Returns to privileged EXEC mode.

Enabling EVPN VXLAN Layer 3 TRM Interworking with MVPN Networks

To configure interworking of Layer 3 TRM with MVPN networks, perform the following steps:

- Configure Layer 3 TRM in the BGP EVPN VXLAN fabric before you enable Layer 3 TRM interworking with MVPN. See [How to Configure Tenant Routed Multicast](#) for detailed steps.
- Configure the MVPN network for the VPNv4 address family. See *Configuring Multicast Virtual Private Network* module of the *IP Multicast Routing Configuration Guide*.



Note Configure MVPN Profile 11 for EVPN VXLAN Layer 3 TRM interworking. For information on how to configure MVPN profiles, refer to [Configure mVPN Profiles Within Cisco IOS](#).


- If internal Border Gateway Protocol (iBGP) is used for peering between the two networks, run the **mdt auto-discovery interworking vxlan-pim** in VRF configuration mode on the border VTEP.

If external Border Gateway Protocol (eBGP) is used for peering between the two networks, run the **mdt auto-discovery interworking vxlan-pim inter-as** in VRF configuration mode on the border VTEP.

Enabling Layer 2 External Connectivity with IEEE 802.1Q Networks

The following image shows a sample topology that illustrates Layer 2 external connectivity with an IEEE 802.1Q network:

Figure 49: Layer 2 External Connectivity with IEEE 802.1Q Networks



You can also connect the EVPN VXLAN network to a firewall in place of the Layer 2 switch in the above image. To configure Layer 2 external connectivity with an IEEE 802.1Q network, perform the following steps on the external Layer 2 switch:


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 <i>interface-id</i> Example: Device(config)# interface GigabitEthernet4/0/1	Enters interface configuration mode for the specified interface. The specified interface must be the interface on the Layer 2 switch through which the EVPN VXLAN network communicates with the IEEE 802.1Q network.
Step 4	switchport mode trunk Example: Device(config-if)# switchport mode trunk	Configures the interface as a trunking VLAN Layer 2 interface.
Step 5	switchport trunk allowed vlan <i>vlan-list</i> Example: Device(config-if)# switchport trunk allowed vlan 201,202	Sets the list of VLANs that are allowed to transmit traffic from this interface in tagged format when the interface is in trunking mode.
Step 6	end Example: Device(config-if)# end	Returns to privileged EXEC mode.

Enabling Layer 2 External Connectivity with a VPLS Network Through an Access VFI

The following illustration shows a single-homed VXLAN network connected to a VPLS over MPLS network through the access VFIs on the border VTEP:

Figure 50: Layer 2 External Connectivity with a VPLS Network Through an Access VFI



Note We recommend you to use Cisco Catalyst 9500 Series - High Performance switches or Cisco Catalyst 9600 Series switches as border VTEPs when you configure Layer 2 external connectivity with a VPLS network.

We recommend you to configure Cisco StackWise Virtual on the border VTEPs in order to achieve physical redundancy when you configure Layer 2 external connectivity with a VPLS network.

Perform the following set of procedures to enable Layer 2 external connectivity with VPLS networks through an access VFI interface:

1. Define the access VFI for the VTEPs.
2. Configure the access VFI as a member of the VLAN on the VTEPs.
3. Configure the EVPN instance as a member of the VLAN on the VTEPs.
4. Configure VPLS on the border VTEP.

Defining an Access VFI on a Border VTEP

To configure an access facing VFI on the VLAN of a border VTEP, perform the following steps:

For more information on configuring VFIs, in the software configuration guide for the required release, go to *Contents* → *Multiprotocol Label Switching (MPLS) Configuration Guide* → *Configuring Virtual Private LAN Service (VPLS) and VPLS BGP-Based Autodiscovery*.

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.

	Command or Action	Purpose
Step 3	l2vpn vfi context <i>vfi-name</i> Example: Device(config) # l2vpn vfi context myVFI	Establishes an Layer 2 VPN VFI between two or more separate networks, and enters VFI configuration mode.
Step 4	vpn id <i>vpn-id</i> Example: Device(config-vfi) # vpn id 1	Configures the VPN ID for the VFI.
Step 5	member ip-address encapsulation mpls Example: Device(config-vfi) # member 10.12.12.5 encapsulation mpls	Specifies the device that forms a point-to-point Layer 2 VPN VFI connection.
Step 6	Repeat step 5 for all devices that form a point-to-point Layer 2 VPN VFI connection.	
Step 7	end Example: Device(config-vfi) # end	Exits VFI configuration mode and enters privileged EXEC mode.

Adding an Access VFI and an EVPN Instance as Members of the VLAN of a Border VTEP

To add an access VFI and an EVPN instance as members of the VLAN of a border 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	vlan configuration <i>vlan-number</i> Example: Device(config) # vlan configuration 11	Enters VLAN feature configuration mode for the specified VLAN interface. Enter the VLAN number that is associated with the Layer 2 VNI configured in the EVPN VXLAN network.
Step 4	member access-vfi <i>vfi-name</i> Example:	Adds the access VFI as a member of the VLAN configuration.

	Command or Action	Purpose
	Device(config-vlan)# member access-vfi myVFI	
Step 5	member evpn-instance <i>evpn-instance-number</i> vni <i>l2-vni-number</i> Example: Device(config-vlan)# member evpn-instance 1 vni 6000	Adds the EVPN instance as a member of the VLAN configuration.
Step 6	end Example: Device(config-vlan)# end	Exits VLAN configuration mode and enters privileged EXEC mode.

Configuring VPLS on a Border VTEP

To configure VPLS on a border VTEP, in the software configuration guide for the required release, see *Contents → Multiprotocol Label Switching (MPLS) Configuration Guide → Configuring Virtual Private LAN Service (VPLS) and VPLS BGP-Based Autodiscovery*.

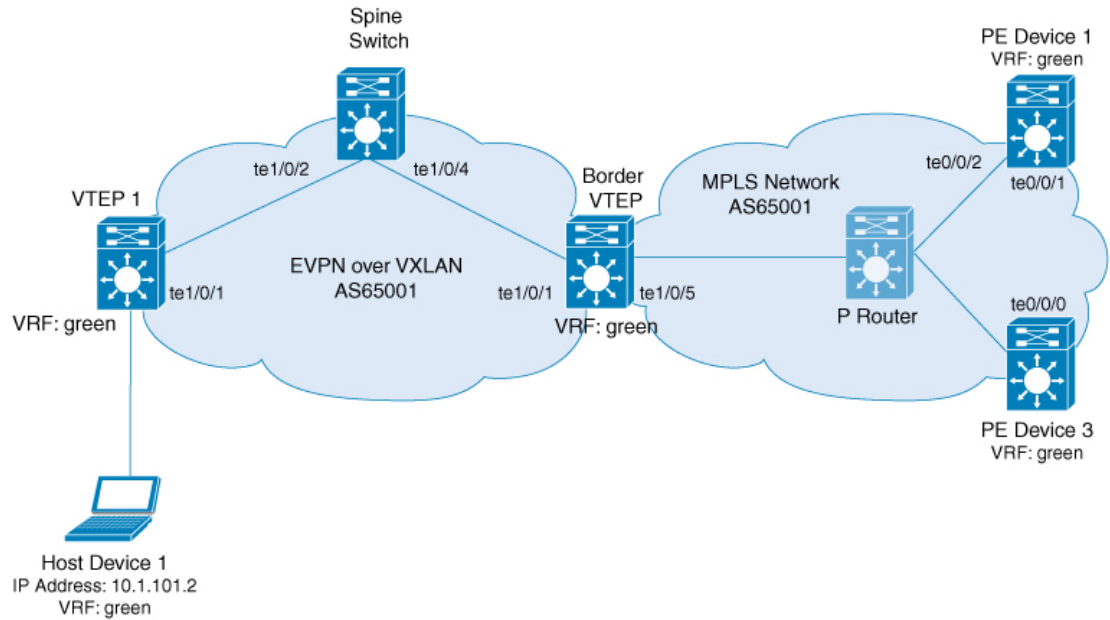
Configuration Examples for EVPN VXLAN External Connectivity

The following section shows the configuration examples for EVPN VXLAN external connectivity to other technologies:

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP

This section provides an example to show how Layer 3 external connectivity with MPLS Layer 3 VPN is enabled for a BGP EVPN VXLAN fabric through iBGP. The example shows how to configure and verify Layer 3 external connectivity with MPLS Layer 3 VPN for the topology shown below:

Figure 51: Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP



The topology shows an EVPN VXLAN network with two VTEPS, VTEP 1 and border VTEP. Border VTEP is connected to an external PE device that belongs to an MPLS network. The BGP EVPN VXLAN fabric and the MPLS network are in the autonomous system number 65001. All the VTEPs, PE devices and, host devices are part of the VRF green. The following tables provide sample configurations for the devices in the topology above.

356585

Table 55: Configuring Spine Switch, Border VTEP and PE Device 1 for Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP

Spine Switch	Border VTEP	PE Device 1
<pre> Spine_switch# show running-config hostname Spine_switch ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ip pim sparse-mode ! 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/4 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 template peer-policy RR-PP route-reflector-client send-community both exit-peer-policy ! template peer-session RR-PS remote-as 65001 update-source Loopback0 exit-peer-session ! bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.4 inherit peer-session RR-PS neighbor 172.16.255.6 inherit peer-session RR-PS ! ! ! </pre>	<pre> Border_VTEP# show running-config hostname Border_VTEP !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 ! mpls label mode all-vrfs protocol all-afs per-vrf ! 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 replication-type ingress ! 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 172.16.255.6 255.255.255.255 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 TenGigabitEthernet1/0/1 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 </pre>	<pre> PE_device_1# show running-config hostname PE_device_1 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 exit-address-family ! interface Loopback0 ip address 172.16.255.101 255.255.255.255 ! interface Loopback1 vrf forwarding green ip address 10.1.255.101 255.255.255.255 ! interface TenGigabitEthernet0/0/1 ip address 172.16.111.101 255.255.255.0 ip router isis cdp enable mpls ip isis network point-to-point ! interface TenGigabitEthernet0/0/2 ip address 172.16.106.101 255.255.255.0 ip router isis negotiation auto cdp enable mpls ip isis network point-to-point ! router isis net 49.0001.1720.1625.5101.00 is-type level-2-only metric-style wide passive-interface Loopback0 ! router bgp 65001 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.103 remote-as 65001 neighbor 172.16.255.103 update-source Loopback0 ! address-family ipv4 exit-address-family ! </pre>

Table 56: Configuring VTEP 1 and PE Device 3 for Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP

VTEP 1	PE Device 3
<pre> VTEP_1# show running-config hostname VTEP_1 ! 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 ! 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 replication-type ingress ! 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 TenGigabitEthernet1/0/1 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 Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 </pre>	<pre> PE_device_3# show running-config hostname PE_device_3 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 exit-address-family ! interface Loopback0 ip address 172.16.255.103 255.255.255.255 ! interface Loopback1 vrf forwarding green ip address 10.1.255.103 255.255.255.255 ! interface TenGigabitEthernet0/0/0 ip address 172.16.111.103 255.255.255.0 ip router isis cdp enable mpls ip isis network point-to-point ! router isis net 49.0001.1720.1625.5103.00 is-type level-2-only metric-style wide passive-interface Loopback0 ! router bgp 65001 template peer-policy RR-PP route-reflector-client send-community both exit-peer-policy ! template peer-session RR-PS remote-as 65001 update-source Loopback0 exit-peer-session ! bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.6 inherit peer-session RR-PS neighbor 172.16.255.101 inherit peer-session RR-PS ! address-family ipv4 exit-address-family ! ! ! ! ! ! </pre>

The following examples provide sample outputs for **show** commands on VTEP 1 and border VTEP to verify external connectivity for the topology configured above:

VTEP 1

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 5 on VTEP 1:

```
VTEP_1# show bgp l2vpn evpn route-type 5 0 10.1.255.103 32
BGP routing table entry for [5][1:1][0][32][10.1.255.103]/17, version 12
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Flag: 0x100
  Not advertised to any peer
  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
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN
Label 0
  Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
  Originator: 172.16.255.103, Cluster list: 172.16.255.1, 172.16.255.6
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F84B914EF38, path: 0x7F84BAFD0E30, pathext: 0x7F84BB42E698
  flags: net: 0x100, path: 0x3, pathext: 0xA1
  Updated on May 20 2020 19:31:08 UTC
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 2 on VTEP 1:

```
VTEP_1# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.2
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.2]/24,
  version 17
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
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
        Router MAC:7C21.0DBD.9548
      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
      net: 0x7F84B914E858, path: 0x7F84BAFD09F8, pathext: 0x7F84BB42E4B8
      flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
      Updated on May 20 2020 19:31:30 UTC
```

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

```
VTEP_1# show ip route vrf green

Routing Table: green
```

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP

```

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, 6 subnets, 2 masks
C       10.1.101.0/24 is directly connected, Vlan101
L       10.1.101.1/32 is directly connected, Vlan101
C       10.1.102.0/24 is directly connected, Vlan102
L       10.1.102.1/32 is directly connected, Vlan102
B       10.1.255.101/32 [200/0] via 172.16.254.6, 00:21:47, Vlan901
B       10.1.255.103/32 [200/0] via 172.16.254.6, 00:21:47, Vlan901

```

Border VTEP

The following example shows the output for the **show mpls ldp neighbor** command on border VTEP:

```

Border_VTEP# show mpls ldp neighbor
Peer LDP Ident: 172.16.111.101:0; Local LDP Ident 172.16.106.6:0
TCP connection: 172.16.111.101.26371 - 172.16.106.6.646
State: Oper; Msgs sent/rcvd: 86/69; Downstream
Up time: 00:32:14
LDP discovery sources:
  TenGigabitEthernet1/0/5, Src IP addr: 172.16.106.101
Addresses bound to peer LDP Ident:
  172.16.111.101 172.16.106.101 172.16.255.101

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 5 on border VTEP:

```

Border_VTEP# show bgp l2vpn evpn route-type 5 0 10.1.255.103 32
BGP routing table entry for [5][1:1][0][32][10.1.255.103]/17, version 7
Paths: (1 available, best #1, table EVPN-BGP-Table)
Flag: 0x100
Advertised to update-groups:
  1
Refresh Epoch 1
Local, (Received from a RR-client), imported path from base
  172.16.255.103 (metric 20) (via default) from 172.16.255.103 (172.16.255.103)
  Origin incomplete, metric 0, localpref 100, valid, internal, best
  EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, local vtep: 172.16.254.6,
VNI Label 50901, MPLS VPN Label 23
  Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
  rx pathid: 0, tx pathid: 0x0
  net: 0x7FED6F808948, path: 0x7FED6D7EDA68, pathext: 0x7FED6D80DE40, exp_net:
0x7FED6F9BF070
  flags: net: 0x100, path: 0x7, pathext: 0xA1
  Updated on May 20 2020 19:22:47 UTC

```

The following example shows the output for the **show bgp vpv4 unicast all** command on border VTEP for the IP address of host device 1:

```
Border_VTEP# show bgp vpv4 unicast all 10.1.101.2
BGP routing table entry for 1:1:10.1.101.2/32, version 10
Paths: (1 available, best #1, table green)
  Advertised to update-groups:
    3
  Refresh Epoch 1
  Local, (Received from a RR-client), imported path from
[2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.2]/24 (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
  Extended Community: RT:1:1 ENCAP:8 Router MAC:7C21.0DBD.9548
  Originator: 172.16.255.4, Cluster list: 172.16.255.1
  Local vxlan vtep:
    vrf:green, vni:50901
    local router mac:0C75.BD67.EF48
    encap:8
    vtep-ip:172.16.254.6
    bdi:Vlan901
  Remote VxLAN:
    Topoid 0x4(vrf green)
    Remote Router MAC:7C21.0DBD.9548
    Encap 8
    Egress VNI 50901
    RTEP 172.16.254.4
  mpls labels in/out IPv4 VRF Aggr:34/nolabel
  rx pathid: 0, tx pathid: 0x0
  Updated on May 20 2020 19:23:11 UTC
```

Spine Switch

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 5 on spine switch:

```
Spine_switch# show bgp l2vpn evpn route-type 5 0 10.1.255.103 32
BGP routing table entry for [5][1:1][0][32][10.1.255.103]/17, version 12
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  Local, (Received from a RR-client)
    172.16.254.6 (metric 2) (via default) from 172.16.255.6 (172.16.255.6)
    Origin incomplete, metric 0, localpref 100, valid, internal, best
    EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN
Label 0
  Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
  Originator: 172.16.255.103, Cluster list: 172.16.255.6
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F54CC99CEF8, path: 0x7F54CC9AD310, pathext: 0x7F54CC9C6998
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 20 2020 19:28:59 UTC
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 2 on spine switch:

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through iBGP

```
Spine_switch# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.2
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.2]/24,
version 14
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  Local, (Received from a RR-client)
    172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      rx pathid: 0, tx pathid: 0x0
      net: 0x7F54CC99CAD8, path: 0x7F54CC9AD088, pathext: 0x7F54CC9C68D8
      flags: net: 0x0, path: 0x3, pathext: 0x81
      Updated on May 20 2020 19:29:22 UTC
```

PE Device 3

The following example shows the output for the **show bgp vpnv4 unicast all** command on PE device 3 for the IP address of host device 1:

```
PE_device_3# show bgp vpnv4 unicast all 10.1.101.2
BGP routing table entry for 1:1:10.1.101.2/32, version 14
Paths: (1 available, best #1, table green)
  Advertised to update-groups:
    3
  Refresh Epoch 1
  Local, (Received from a RR-client)
    172.16.255.6 (metric 20) (via default) from 172.16.255.6 (172.16.255.6)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:1:1 ENCAP:8 Router MAC:7C21.0DBD.9548
      Originator: 172.16.255.4, Cluster list: 172.16.255.6, 172.16.255.1
      mpls labels in/out nolabel/34
      rx pathid: 0, tx pathid: 0x0
      Updated on May 20 2020 11:27:25 UTC
```

The following example shows the output for the **show ip route vrf green** command on PE device 3:

```
PE_device_3# 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, 7 subnets, 2 masks
```



```


B      10.1.101.0/24 [200/0] via 172.16.255.6, 00:28:12
B      10.1.101.1/32 [200/0] via 172.16.255.6, 00:28:10
B      10.1.101.2/32 [200/0] via 172.16.255.6, 00:27:48
B      10.1.102.0/24 [200/0] via 172.16.255.6, 00:28:12
B      10.1.102.1/32 [200/0] via 172.16.255.6, 00:28:10
B      10.1.255.101/32 [200/0] via 172.16.255.101, 00:28:09
C      10.1.255.103/32 is directly connected, Loopback1

```

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP

This section provides an example to show how Layer 3 external connectivity with MPLS Layer 3 VPN is enabled for a BGP EVPN VXLAN fabric through eBGP. The example shows how to configure and verify Layer 3 external connectivity with MPLS Layer 3 VPN for the topology shown below:

Figure 52: Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP



The topology shows an EVPN VXLAN network with two VTEPs, VTEP 1 and border VTEP. Border VTEP is connected to an external PE device that belongs to an MPLS network. The BGP EVPN VXLAN fabric is in the autonomous system number 65001. The MPLS network is in the autonomous system number 65002. All the VTEPs, PE devices, and host devices are part of the VRF green. The following tables provide sample configurations for the devices in the topology above.

Table 57: Configuring Spine Switch, Border VTEP and PE Device 1 for Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP

Spine Switch	Border VTEP	PE Device 1
<pre> Spine_switch# show running-config hostname Spine_switch ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ip pim sparse-mode ! 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 TenGigabitEthernet1/0/2 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 TenGigabitEthernet1/0/4 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 template peer-policy RR-PP route-reflector-client send-community both exit-peer-policy ! template peer-session RR-PS remote-as 65001 update-source Loopback0 exit-peer-session ! bgp router-id 172.16.255.1 bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.4 inherit peer-session RR-PS neighbor 172.16.255.6 inherit peer-session RR-PS ! address-family ipv4 exit-address-family ! </pre>	<pre> Border_VTEP# show running-config hostname Border_VTEP ! 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 ! mpls label mode all-vrfs protocol all-afs per-vrf ! 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 replication-type ingress ! 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 172.16.255.6 255.255.255.255 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 TenGigabitEthernet1/0/1 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 </pre>	<pre> PE_device_1# show running-config hostname PE_device_1 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 exit-address-family ! interface Loopback0 ip address 172.16.255.101 255.255.255.255 ! interface Loopback1 vrf forwarding green ip address 10.1.255.101 255.255.255.255 ! interface TenGigabitEthernet0/0/1 ip address 172.16.111.101 255.255.255.0 ip router isis cdp enable mpls ip isis network point-to-point ! interface TenGigabitEthernet0/0/2 ip address 172.16.106.101 255.255.255.0 negotiation auto cdp enable mpls bgp forwarding ! router isis net 49.0001.1720.1625.5101.00 is-type level-2-only metric-style wide passive-interface Loopback0 ! router bgp 65002 bgp log-neighbor-changes no bgp default ipv4-unicas no bgp default route-target filter neighbor 172.16.106.6 remote-as 65001 neighbor 172.16.255.6 remote-as 65001 neighbor 172.16.255.6 ebgp-multihop 255 neighbor 172.16.255.6 update-source Loopback0 neighbor 172.16.255.103 remote-as 65002 neighbor 172.16.255.103 update-source Loopback0 </pre>

Table 58: Configuring VTEP 1 and PE Device 3 for Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP

VTEP 1	PE Device 3
<pre>VTEP_1# show running-config hostname VTEP_1! ! 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 ! 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 replication-type ingress ! 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 TenGigabitEthernet1/0/1 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 !</pre>	<pre>PE_device_3# show running-config hostname PE_device_3 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 exit-address-family ! interface Loopback0 ip address 172.16.255.103 255.255.255.255 ! interface Loopback1 vrf forwarding green ip address 10.1.255.103 255.255.255.255 ! interface TenGigabitEthernet0/0/0 ip address 172.16.111.103 255.255.255.0 ip router isis cdp enable mpls ip isis network point-to-point ! router isis net 49.0001.1720.1625.5103.00 is-type level-2-only metric-style wide passive-interface Loopback0 ! router bgp 65002 template peer-policy RR-PP route-reflector-client send-community both exit-peer-policy ! template peer-session RR-PS remote-as 65002 update-source Loopback0 exit-peer-session ! bgp log-neighbor-changes no bgp default ipv4-unicast neighbor 172.16.255.101 inherit peer-session RR-PS ! address-family ipv4 exit-address-family ! ! !</pre>

The following examples provide sample outputs for **show** commands on the devices to verify external connectivity for the topology configured above:

VTEP 1

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 5 on VTEP 1:

```
VTEP_1# show bgp l2vpn evpn route-type 5 0 10.1.255.103 32
BGP routing table entry for [5][1:1][0][32][10.1.255.103]/17, version 36
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 1
  65002
    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
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN
Label 0
  Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
  Originator: 172.16.255.6, Cluster list: 172.16.255.1
  rx pathid: 0, tx pathid: 0x0
  net: 0x7F84BB35A5C8, path: 0x7F84B913E010, pathext: 0x7F84BB54A8A8
  flags: net: 0x0, path: 0x3, pathext: 0x81
  Updated on May 21 2020 13:56:28 UTC
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 2 on VTEP 1:

```
VTEP_1# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.2
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.2]/24,
version 37
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
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
      Router MAC:7C21.0DBD.9548
      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
      net: 0x7F84BB35A468, path: 0x7F84B913DF38, pathext: 0x7F84BB54A848
      flags: net: 0x0, path: 0x4000028000003, pathext: 0x81
      Updated on May 21 2020 14:00:49 UTC
```

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

```
VTEP_1# 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
```

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP

```

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, 6 subnets, 2 masks
C    10.1.101.0/24 is directly connected, Vlan101
L    10.1.101.1/32 is directly connected, Vlan101
C    10.1.102.0/24 is directly connected, Vlan102
L    10.1.102.1/32 is directly connected, Vlan102
B    10.1.255.101/32 [200/0] via 172.16.254.6, 00:06:25, Vlan901
B    10.1.255.103/32 [200/0] via 172.16.254.6, 00:05:54, Vlan901

```

Border VTEP

The following example shows the output for the **show bgp vpnv4 unicast all** command on border VTEP for the IP address of the external device:

```

Border_VTEP# show bgp vpnv4 uni all 10.1.255.103/32
BGP routing table entry for 1:1:10.1.255.103/32, version 9
Paths: (1 available, best #1, table green)
  Not advertised to any peer
  Refresh Epoch 1
  65002
    172.16.255.101 (via default) from 172.16.255.101 (172.16.255.101)
      Origin incomplete, localpref 100, valid, external, best
      Extended Community: RT:1:1
      Local vxlan vtep:
        vrf:green, vni:50901
        local router mac:0C75.BD67.EF48
        encaps:8
        vtep-ip:172.16.254.6
        bdi:Vlan901
      mpls labels in/out nolabel/16
      rx pathid: 0, tx pathid: 0x0
      Updated on May 21 2020 13:48:09 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 5 on border VTEP:

```

Border_VTEP# show bgp l2vpn evpn route-type 5 0 10.1.255.103 32
BGP routing table entry for [5][1:1][0][32][10.1.255.103]/17, version 32
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  65002, imported path from base
    172.16.255.101 (via default) from 172.16.255.101 (172.16.255.101)
      Origin incomplete, localpref 100, valid, external, best
      EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, local vtep: 172.16.254.6,
VNI Label 50901, MPLS VPN Label 16
      Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
      rx pathid: 0, tx pathid: 0x0
      net: 0x7FED704944D0, path: 0x7FED704A4CA0, pathext: 0x7FED6DA6E250, exp_net:
0x7FED6F812678

```



```
flags: net: 0x0, path: 0x7, pathext: 0x81
Updated on May 21 2020 13:48:09 UTC
```

The following example shows the output for the **show mpls forwarding-table** command on border VTEP:

```
Border_VTEP# show mpls forwarding-table
Local   Outgoing Prefix          Bytes Label   Outgoing  Next Hop
Label   Label   or Tunnel Id    Switched      interface
16      No Label IPv4 VRF[V]      156          aggregate/green
17      Pop Label 172.16.106.101/32 \
                                     228          Te1/0/5    172.16.106.101
18      Pop Label 172.16.255.101/32 \
                                     0            Te1/0/5    172.16.106.101
```

The following example shows the output for the **show bgp vpnv4 unicast all** command on border VTEP for the IP address of host device 1:

```
Border_VTEP# show bgp vpnv4 uni all 10.1.101.2/32
BGP routing table entry for 1:1:10.1.101.2/32, version 10
Paths: (1 available, best #1, table green)
  Advertised to update-groups:
    1
  Refresh Epoch 4
  Local, imported path from [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.2]/24
(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
  Extended Community: RT:1:1 ENCAP:8 Router MAC:7C21.0DBD.9548
  Originator: 172.16.255.4, Cluster list: 172.16.255.1
  Local vxlan vtep:
    vrf:green, vni:50901
    local router mac:0C75.BD67.EF48
    encap:8
    vtep-ip:172.16.254.6
    bdi:Vlan901
  Remote VxLAN:
    Topoid 0x9(vrf green)
    Remote Router MAC:7C21.0DBD.9548
    Encap 8
    Egress VNI 50901
    RTEP 172.16.254.4
  mpls labels in/out IPv4 VRF Aggr:16/nolabel
  rx pathid: 0, tx pathid: 0x0
  Updated on May 21 2020 13:52:30 UTC
```

Spine Switch

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 5 on spine switch:

```
Spine_switch# show bgp l2vpn evpn route-type 5 0 10.1.255.103 32
BGP routing table entry for [5][1:1][0][32][10.1.255.103]/17, version 23
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  65002, (Received from a RR-client)
    172.16.254.6 (metric 2) (via default) from 172.16.255.6 (172.16.255.6)
```

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP

```

Origin incomplete, metric 0, localpref 100, valid, internal, best
EVPN ESI: 00000000000000000000, Gateway Address: 0.0.0.0, VNI Label 50901, MPLS VPN
Label 0
Extended Community: RT:1:1 ENCAP:8 Router MAC:0C75.BD67.EF48
rx pathid: 0, tx pathid: 0x0
net: 0x7F54CC95FAB8, path: 0x7F54CCA542F8, pathext: 0x7F54CC9707B0
flags: net: 0x0, path: 0x3, pathext: 0x81
Updated on May 21 2020 13:54:20 UTC

```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 2 on spine switch:

```

Spine_switch# show bgp l2vpn evpn route-type 2 0 44d3ca286cc1 10.1.101.2
BGP routing table entry for [2][172.16.254.4:101][0][48][44D3CA286CC1][32][10.1.101.2]/24,
version 24
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  Local, (Received from a RR-client)
    172.16.254.4 (metric 2) (via default) from 172.16.255.4 (172.16.255.4)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      EVPN ESI: 00000000000000000000, Label1 10101, Label2 50901
      Extended Community: RT:1:1 RT:65001:101 ENCAP:8
        Router MAC:7C21.0DBD.9548
        rx pathid: 0, tx pathid: 0x0
        net: 0x7F54CC95F958, path: 0x7F54CCA54220, pathext: 0x7F54CC970750
        flags: net: 0x0, path: 0x3, pathext: 0x81
        Updated on May 21 2020 13:58:41 UTC

```

PE Device 1

The following example shows the output for the **show bgp vpnv4 unicast all** command on PE device 1 for the IP address of host device 1:

```

PE_device_1# show bgp vpnv4 unicast all 10.1.255.103/32
BGP routing table entry for 1:1:10.1.101.2/32, version 14
Paths: (1 available, best #1, table green)
  Advertised to update-groups:
    1
  Refresh Epoch 1
  65001
    172.16.255.6 (via default) from 172.16.255.6 (172.16.255.6)
      Origin incomplete, localpref 100, valid, external, best
      Extended Community: RT:1:1 ENCAP:8 Router MAC:7C21.0DBD.9548
      mpls labels in/out 22/16
      rx pathid: 0, tx pathid: 0x0
      Updated on May 21 2020 05:57:06 UTC

```

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

```

PE_device_1# 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, 7 subnets, 2 masks
B    10.1.101.0/24 [20/0] via 172.16.255.6, 00:28:09
B    10.1.101.1/32 [20/0] via 172.16.255.6, 00:28:09
B    10.1.101.2/32 [20/0] via 172.16.255.6, 00:23:17
B    10.1.102.0/24 [20/0] via 172.16.255.6, 00:28:09
B    10.1.102.1/32 [20/0] via 172.16.255.6, 00:28:09
C    10.1.255.101/32 is directly connected, Loopback1
B    10.1.255.103/32 [200/0] via 172.16.255.103, 00:28:09

```

PE Device 3

The following example shows the output for the **show bgp vpnv4 unicast all** command on PE device 3 for the IP address of host device 1:

```

PE_device_3# show bgp vpnv4 unicast all 10.1.101.2/32
BGP routing table entry for 1:1:10.1.101.2/32, version 14
Paths: (1 available, best #1, table green)
  Not advertised to any peer
  Refresh Epoch 1
  65001, (Received from a RR-client)
    172.16.255.101 (metric 10) (via default) from 172.16.255.101 (172.16.255.101)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: RT:1:1 ENCAP:8 Router MAC:7C21.0DBD.9548
      mpls labels in/out nolabel/22
      rx pathid: 0, tx pathid: 0x0
      Updated on May 21 2020 05:56:46 UTC

```

The following example shows the output for the **show ip route vrf** command on PE device 3:

```

PE_device_3# 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, 7 subnets, 2 masks

```

Example: Enabling Layer 3 External Connectivity with MPLS Layer 3 VPN through eBGP

```
B      10.1.101.0/24 [200/0] via 172.16.255.101, 00:29:09
B      10.1.101.1/32 [200/0] via 172.16.255.101, 00:29:09
B      10.1.101.2/32 [200/0] via 172.16.255.101, 00:24:17
B      10.1.102.0/24 [200/0] via 172.16.255.101, 00:29:09
B      10.1.102.1/32 [200/0] via 172.16.255.101, 00:29:09
B      10.1.255.101/32 [200/0] via 172.16.255.101, 00:29:09
C      10.1.255.103/32 is directly connected, Loopback1
```



CHAPTER 12

Configuring Multi-Homing in a BGP EVPN VXLAN Fabric

- [Restrictions for Multi-Homing in a BGP EVPN VXLAN Fabric, on page 647](#)
- [Information About Multi-Homing in a BGP EVPN VXLAN Fabric, on page 647](#)
- [How to Configure Multi-Homing in a BGP EVPN VXLAN Fabric, on page 653](#)
- [Configuration Examples for Multi-Homing in a BGP EVPN VXLAN Fabric, on page 658](#)

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 Catalyst 9000 Series switch 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 Catalyst 9000 Series 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 53: Single-Homed Topology



357247

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:

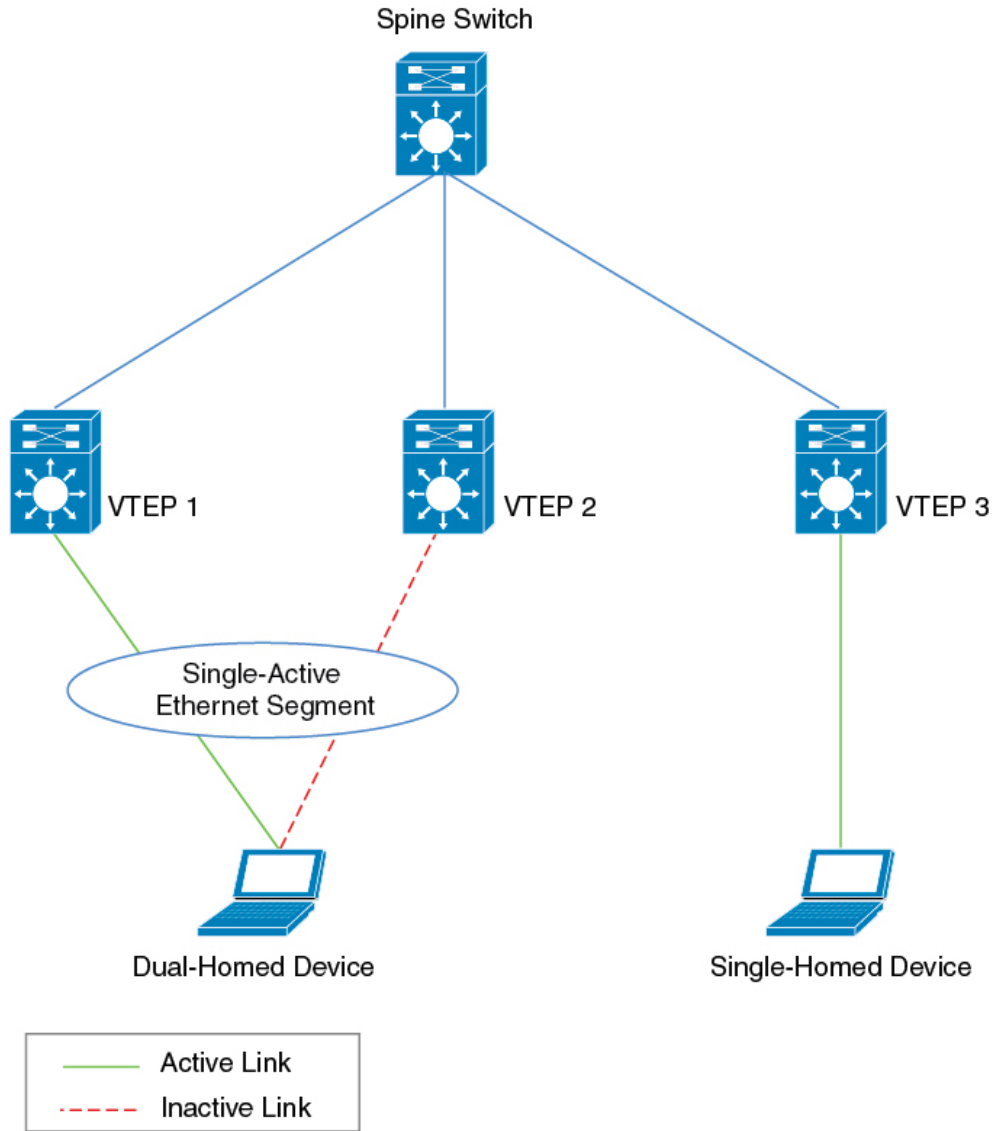
- Dual-homed device
- Dual-homed network

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 54: Dual-Homed Device Topology



957245

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.

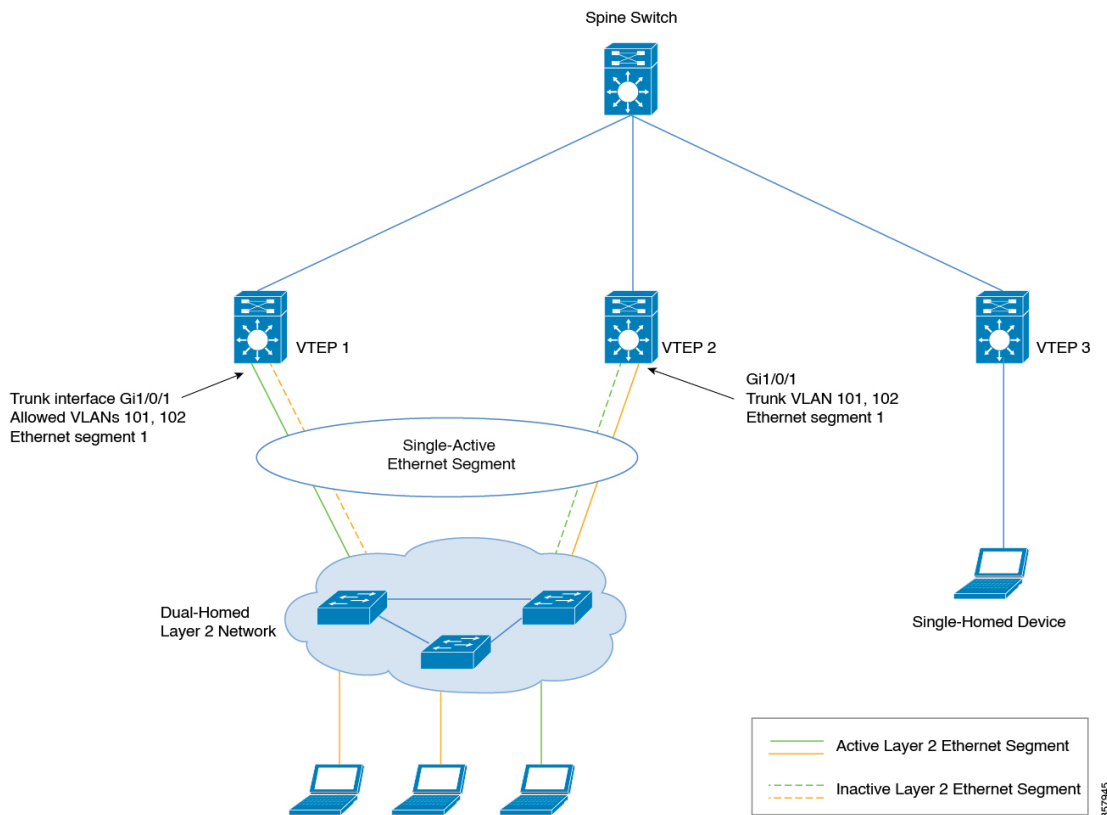
A dual-homed network topology results in a situation where the network splits into two different networks if the connectivity between the host or access devices is lost. To avoid this scenario, redundancy must also be enabled within the dual-homed network.

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 55: 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.

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 655](#).

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 657](#).

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 101](#) 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.

	Command or Action	Purpose
Step 3	l2vpn evpn ethernet-segment <i>ethernet-segment-id</i> 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-address } Example: Device(config-evpn-es)# identifier type 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: <ul style="list-style-type: none"> • 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 <i>redundancy-type</i> Example: Device(config-evpn-es)# redundancy single-active	Configures the redundancy type for the ethernet segment.
Step 6	df-election wait-time <i>time-period</i> 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:	Enters global configuration mode.

	Command or Action	Purpose
	Device# <code>configure terminal</code>	
Step 3	interface <i>interface-id</i> Example: Device(config)# interface GigabitEthernet1/0/10	Specifies the interface, and enters interface configuration mode.
Step 4	evpn ethernet-segment <i>ethernet-segment-id</i> 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 56: 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 57: 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 58: 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 664](#) section.

Table 59: Configuring Dual-Homing with Single-Active Redundancy using VTEP 2 and VTEP 3

VTEP 2	VTEP 3
<pre> 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 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 ! </pre>	<pre> 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 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 ! </pre>

VTEP 2	VTEP 3
<pre> interface Loopback1 ip address 172.16.254.4 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.182 255.255.255.0 negotiation auto ! 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 </pre>	<pre> interface Loopback1 ip address 172.16.254.5 255.255.255.255 ip ospf 1 area 0 ! interface GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.183 255.255.255.0 negotiation auto ! interface GigabitEthernet1/0/1 no switchport ip address 172.16.15.5 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/2 no switchport ip address 172.16.25.5 255.255.255.0 ip ospf network point-to-point ip ospf 1 area 0 ! interface GigabitEthernet1/0/10 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 ! 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 </pre>

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

VTEP 2	VTEP 3
<pre> ! 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 connected redistribute static exit-address-family ! end ! Leaf-02# </pre>	<pre> ! 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 ! end ! Leaf-03# </pre>

Table 60: 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
<pre> 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 GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.180 255.255.255.0 negotiation auto ! 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 ! </pre>	<pre> 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 GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.180 255.255.255.0 negotiation auto ! 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 ! </pre>	<pre> 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 ! </pre>

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

Spine Switch 1	Spine Switch 2	VTEP 1
<pre> 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# </pre>	<pre> 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# </pre>	<pre> interface GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.179 255.255.255.0 negotiation auto ! 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 ! </pre>

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 665](#)
- [Outputs to Verify the Configuration on VTEP 2, on page 667](#)
- [Outputs to Verify the Configuration on VTEP 3, on page 670](#)

- [Outputs to Verify the Configuration on Spine Switch 1, on page 674](#)
- [Outputs to Verify the Configuration on Spine Switch 2, on page 676](#)

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:

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

```
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      V1101:0      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      Gi1/0/10:101 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 664](#).

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
```

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

```
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	?
*>	::			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	?
*>	::			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:

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

```
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	V1101:0	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 664.

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:          Gil/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
```

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

```
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      Gi1/0/10:101      0
101      0     L2VPN 44d3.ca28.6cc2      Gi1/0/10:101      0
101      0     L2VPN 7c21.0dbd.2741      V1101:0           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      Gi1/0/10:101      0
101      0     L2VPN 44d3.ca28.6cc2      Gi1/0/10:101      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):            Gi1/0/10:101
EVPN Instance:          101

```

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

```

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):          Gi1/0/10:101

```

```
Leaf-03#
```

Return to [Verifying Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric](#), on page 664.

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:

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

```

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 664](#).

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 [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	?
* 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	?
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 [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24	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 [2][172.16.254.3:101][0][48][F4CFE24334C1][0][]/20	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 [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24	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 [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24	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 [2][172.16.254.5:101][0][48][44D3CA286C82][0][]/20	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 [2][172.16.254.5:101][0][48][44D3CA286CC2][0][]/20	172.16.254.5	0	100	0	?

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

```

* 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

```

```

O      [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
      172.16.254.4/32
O      [110/2] via 172.16.24.4, 03:37:28, GigabitEthernet1/0/2
      172.16.254.5/32
O      [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
      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
      172.16.255.4/32
O      [110/2] via 172.16.24.4, 03:37:34, GigabitEthernet1/0/2
      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 664.

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



CHAPTER 13

Configuring Private VLANs in a BGP EVPN VXLAN Fabric

- [Restrictions for Private VLANs in a BGP EVPN VXLAN Fabric, on page 681](#)
- [Information About Private VLANs in a BGP EVPN VXLAN Fabric, on page 681](#)
- [How to Configure Private VLANs in a BGP EVPN VXLAN Fabric, on page 687](#)
- [Configuration Examples for Private VLANs in a BGP EVPN VXLAN Fabric, on page 691](#)

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.

For more information about PVLANS and the steps to configure PVLANS, see "Configuring Private VLANs" module in the *VLAN Configuration Guide* for the applicable release.

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 on 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 59: PVLAN Extension in a BGP EVPN VXLAN Fabric



Traffic Forwarding for Private VLANs in a BGP EVPN VXLAN Fabric

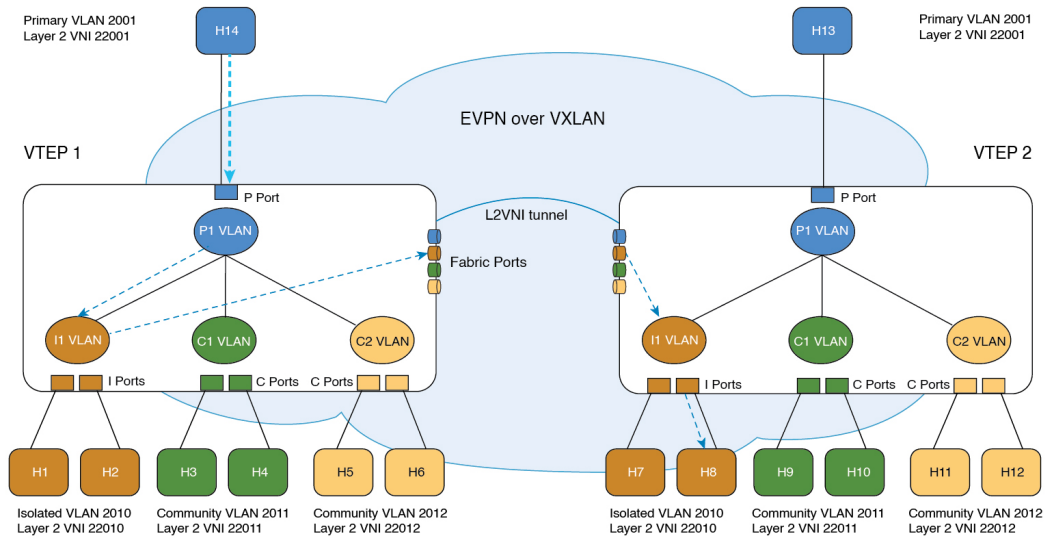
You can forward known unicast and broadcast, unknown unicast, and multicast (BUM) traffic between PVLANs in a BGP EVPN VXLAN fabric. On the Source VTEP, the forwarding process on the access PVLAN ports (promiscuous, isolated, community) adheres to the baseline PVLAN forwarding. With BGP EVPN VXLAN enabled 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 60: Unicast Traffic from Promiscuous Port : H14 to H8




357621

Figure 61: Unicast Traffic from Isolated Port: H1 to H13



357622

Figure 62: Unicast Traffic from Community Port: H6 to H7



357623

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 [Configuring EVPN VXLAN Layer 2 Overlay Network](#)). 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 63: BUM Traffic from Promiscuous Port (H14)




Figure 64: BUM Traffic from Isolated Port (H1)





Figure 65: BUM Traffic from Community Port (H5)



357620

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 95](#)). On the destination VTEP, traffic is routed from the core 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 <i>vlan-id</i> 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.

	Command or Action	Purpose
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}	Associates a PVLAN host port or maps a PVLAN promiscuous port to a primary VLAN.

	Command or Action	Purpose
	Example: <pre>Device(config-if) # switchport private-vlan host-association 101 104</pre>	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: <pre>Device(config-if) # end</pre>	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: <pre>Device> enable</pre>	Enters privileged EXEC mode. Enter password, if prompted.
Step 2	configure terminal Example: <pre>Device# configure terminal</pre>	Enters global configuration mode.
Step 3	vlan configuration <i>vlan-id</i> Example: <pre>Device(config)# vlan configuration 101</pre>	Enters VLAN configuration mode for the specified PVLAN interface.
Step 4	member evpn-instance <i>evpn-instance-id</i> vni <i>layer2-vni-id</i>	Adds EVPN instance as a member of the PVLAN configuration.

	Command or Action	Purpose
	Example: Device (config-vlan) # member evpn-instance 1 vni 6000	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 66: 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 61: Configuring VTEP 1, VTEP 2, and VTEP 3 for PVLAN Extension in a BGP EVPN VXLAN Fabric

VTEP 1	VTEP 2	VTEP 3
<pre>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</pre>	<pre>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</pre>	<pre>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</pre>

VTEP 1	VTEP 2	VTEP 3
<pre> ! l2vpn evpn instance 203 vlan-based encapsulation vxlan ! l2vpn evpn instance 204 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 103 member evpn-instance 103 vni 10103 vlan configuration 104 member evpn-instance 104 vni 10104 vlan configuration 201 member evpn-instance 201 vni 10201 vlan configuration 202 member evpn-instance 202 vni 10202 vlan configuration 203 member evpn-instance 203 vni 10203 vlan configuration 204 member evpn-instance 204 vni 10204 vlan configuration 901 member vni 50901 ! vlan 101 private-vlan primary private-vlan association 102-104 ! vlan 102 private-vlan community ! vlan 103 private-vlan community ! vlan 104 private-vlan isolated ! vlan 201 private-vlan primary private-vlan association 202-204 ! vlan 202 private-vlan community ! </pre>	<pre> ! l2vpn evpn instance 203 vlan-based encapsulation vxlan ! l2vpn evpn instance 204 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 103 member evpn-instance 103 vni 10103 vlan configuration 104 member evpn-instance 104 vni 10104 vlan configuration 201 member evpn-instance 201 vni 10201 vlan configuration 202 member evpn-instance 202 vni 10202 vlan configuration 203 member evpn-instance 203 vni 10203 vlan configuration 204 member evpn-instance 204 vni 10204 vlan configuration 901 member vni 50901 ! vlan 101 private-vlan primary private-vlan association 102-104 ! vlan 102 private-vlan community ! vlan 103 private-vlan community ! vlan 104 private-vlan isolated ! vlan 201 private-vlan primary private-vlan association 202-204 ! vlan 202 private-vlan community ! </pre>	<pre> ! l2vpn evpn instance 203 vlan-based encapsulation vxlan ! l2vpn evpn instance 204 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 103 member evpn-instance 103 vni 10103 vlan configuration 104 member evpn-instance 104 vni 10104 vlan configuration 201 member evpn-instance 201 vni 10201 vlan configuration 202 member evpn-instance 202 vni 10202 vlan configuration 203 member evpn-instance 203 vni 10203 vlan configuration 204 member evpn-instance 204 vni 10204 vlan configuration 901 member vni 50901 ! vlan 101 private-vlan primary private-vlan association 102-104 ! vlan 102 private-vlan community ! vlan 103 private-vlan community ! vlan 104 private-vlan isolated ! vlan 201 private-vlan primary private-vlan association 202-204 ! vlan 202 private-vlan community ! </pre>

VTEP 1	VTEP 2	VTEP 3
<pre> 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 ! </pre>	<pre> 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 ! </pre>	<pre> 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 GigabitEthernet0/0 vrf forwarding Mgmt-vrf ip address 10.62.149.183 255.255.255.0 negotiation auto ! 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 ! </pre>

VTEP 1	VTEP 2	VTEP 3
<pre> 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 ! </pre>	<pre> 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 ! </pre>	<pre> 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 ! </pre>

VTEP 1	VTEP 2	VTEP 3
<pre> 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# </pre>	<pre> 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# </pre>	<pre> 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# </pre>

Table 62: Configuring Spine Switch 1 and Spine Switch 2 for PVLAN Extension in a BGP EVPN VXLAN Fabric

Spine Switch 1	Spine Switch 2
<pre> 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 ! </pre>	<pre> 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 ! </pre>

Spine Switch 1	Spine Switch 2
<pre> 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# </pre>	<pre> 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# </pre>

Verifying PVLAN Extension in a BGP EVPN VXLAN Fabric

The following sections provide sample outputs for **show** commands to verify the PVLAN extension on the devices in the topology configured above:

- [#unique_231 unique_231_Connect_42_section_ad2_bfg_dqb](#)
- [#unique_231 unique_231_Connect_42_section_b4h_bfg_dqb](#)
- [#unique_231 unique_231_Connect_42_section_p1l_bfg_dqb](#)
- [#unique_231 unique_231_Connect_42_section_cl4_bfg_dqb](#)
- [#unique_231 unique_231_Connect_42_section_mbs_bfg_dqb](#)

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          Gi1/0/3:101
f4cf.e243.34c3  101    101   0000.0000.0000.0000.0000  0          Gi1/0/4:101
f4cf.e243.34c4  101    101   0000.0000.0000.0000.0000  0          Gi1/0/5:101
f4cf.e243.34c2  102    102   0000.0000.0000.0000.0000  0          Gi1/0/3:102
f4cf.e243.34c3  103    103   0000.0000.0000.0000.0000  0          Gi1/0/4:103
f4cf.e243.34c4  104    104   0000.0000.0000.0000.0000  0          Gi1/0/5:104
```

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  V1101:0
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  Gi1/0/5:101
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 ?
*>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_RT  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          Gi1/0/11:101
44d3.ca28.6cc4  101     101   0000.0000.0000.0000.0000  0          Gi1/0/12:101
44d3.ca28.6cc5  101     101   0000.0000.0000.0000.0000  0          Gi1/0/13:101
44d3.ca28.6cc3  102     102   0000.0000.0000.0000.0000  0          Gi1/0/11:102
44d3.ca28.6cc4  103     103   0000.0000.0000.0000.0000  0          Gi1/0/12:103
44d3.ca28.6cc5  104     104   0000.0000.0000.0000.0000  0          Gi1/0/13:104

```

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
101      0     L2VPN 44d3.ca28.6cc4                           10.1.101.14                          Gi1/0/12:101
101      0     L2VPN 44d3.ca28.6cc5                           10.1.101.15                          Gi1/0/13:101
101      0     L2VPN 7c21.0dbd.9541                           10.1.101.1                            V1101:0
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
*>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          Gi1/0/16:201
44d3.ca28.6cc7  201    201   0000.0000.0000.0000.0000  0          Gi1/0/17:201
44d3.ca28.6cc8  201    201   0000.0000.0000.0000.0000  0          Gi1/0/18:201
44d3.ca28.6cc6  202    202   0000.0000.0000.0000.0000  0          Gi1/0/16:202
44d3.ca28.6cc7  203    203   0000.0000.0000.0000.0000  0          Gi1/0/17:203
44d3.ca28.6cc8  204    204   0000.0000.0000.0000.0000  0          Gi1/0/18:204
```

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
201    0     L2VPN 44d3.ca28.6cc7                    10.1.102.4      Gi1/0/17:201
201    0     L2VPN 44d3.ca28.6cc8                    10.1.102.5      Gi1/0/18:201
201    0     L2VPN 7c21.0dbd.274c                    10.1.201.1      V1201:0
```

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					
>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 [5][1:1][0][24][10.1.201.0]/17
      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 BGP EVPN VXLAN over IPsec

This chapter describes how to configure BGP EVPN VXLAN over IPsec. VXLAN support over IPsec provides an overlay which seamlessly extends secured services to enterprise branch locations.

- [Restrictions for BGP EVPN VXLAN over IPsec, on page 715](#)
- [Information about BGP EVPN VXLAN over IPsec, on page 715](#)
- [Workflow to Configure BGP EVPN VXLAN over IPsec, on page 716](#)
- [Configuration Example for BGP EVPN VXLAN over IPsec, on page 718](#)

Restrictions for BGP EVPN VXLAN over IPsec

- BGP EVPN VXLAN over IPsec is supported only on the Cisco Catalyst 9300X Series switch.
- Tenant Routed Multicast over BGP EVPN VXLAN over IPsec tunnel is currently not supported.
- A Catalyst 9300X switch supports a maximum of 128 IPsec tunnels. Hence the size of the BGP EVPN VXLAN fabric over IPsec tunnels is limited to 128 EVPN VXLAN VTEPs terminating the IPsec tunnels.

Information about BGP EVPN VXLAN over IPsec

Internet Protocol Security (IPsec) is a framework of open standards that are developed by the IETF. IPsec provides security for the transmission of sensitive information over unprotected networks such as the Internet. IPsec acts at the network layer, protecting and authenticating IP packets between participating IPsec devices such as hosts, and security gateways. For more information, see "Configuring IPsec" section in the *Security Configuration Guide*.

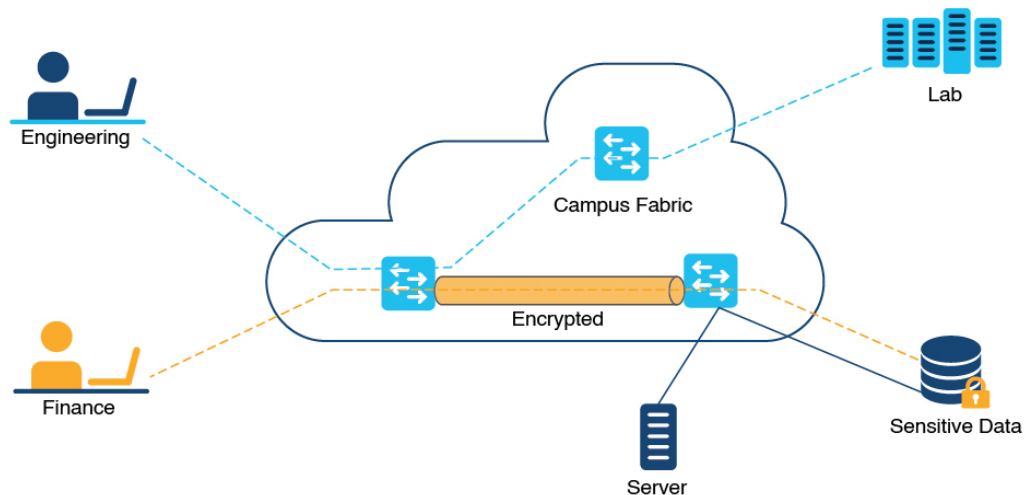
When a BGP EVPN VXLAN network is connected to an external network, the VXLAN traffic flows over the public network or internet, which is unencrypted and prone to data compromise. The traditional GRE Tunnel solution requires a unique tunnel and routing protocol session for each virtual routing and forwarding (VRF) or Layer 3 virtual network instance (L3VNI). The BGP EVPN VXLAN control plane simplifies the overlay network by having a single BGP session with the remote designated IPsec peer to enable large number of VRFs. BGP EVPN VXLAN, when deployed with IPsec, reduces operational complexity and enables secure remote branch access over WAN or internal zero-trust LAN network environments.

BGP EVPN VXLAN over IPsec enables secure encrypted network virtualization with Cisco Catalyst 9300X-based crypto hardware acceleration.

Zero-trust LAN network environments

A campus LAN network with Cisco Catalyst 9300X in the access layer can build secure, encrypted BGP EVPN VXLAN fabric to support a zero-trust network environment. A Cisco Catalyst 9300X at the access layer establishes IPsec tunnel with a Cisco Catalyst 9300X spine border that supports the BGP Route-Reflector functionality and external connectivity.

Figure 67: Campus LAN With Secure BGP EVPN VXLAN Network



Remote Access

Cisco Catalyst 9300X can be deployed as a border VTEP at a branch site to provide secure connectivity to the campus network over a WAN, with IPsec encryption.

Workflow to Configure BGP EVPN VXLAN over IPsec

Before you begin

Ensure that the devices have the correct license to run IPsec and EVPN VXLAN. For license information, refer [Cisco 9300 Series Switch Data Sheet](#).

Procedure

Step 1

Configure the BGP EVPN VXLAN overlay:

Based on the network reachability and segmentation requirements, configure the appropriate overlay topology. Refer to the respective sections for more details on each of the following overlay network segmentations:

- a) Configure L2 overlay: Perform all the configuration tasks that are listed in [How to Configure EVPN VXLAN Layer 2 Overlay Network](#).
- b) Configure L3 overlay: Perform all the configuration tasks that are listed in [How to Configure EVPN VXLAN Layer 3 Overlay Network](#).

- c) Configure Distributed Anycast Gateway (DAG) or Centralized Gateway (CGW): Perform all the tasks that are listed in [Configuring EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast Gateway](#).

Step 2 Configure IPsec in the underlay: Perform all the tasks that are listed in the "Configuring IPsec" section of the *Security Configuration Guide*.

Note Each secure packet that is transmitted has an IPsec header in addition to the BGP EVPN VXLAN header. Adjust the System MTU and TCP MSS size accordingly.

Step 3 Configure BGP neighborship over IPsec tunnel.

- a) To establish BGP neighborship and for loopback interface reachability, configure interior gateway protocol (IGP) over IPsec tunnel.

Here is a sample configuration snippet:

NVE loopback and IPsec tunnel in one IGP instance: the following commands establish NVE or VXLAN neighborship over IPsec tunnel:

```
interface Loopback1
  description NVE Loopback
  ip address 172.16.254.1 255.255.255.255
  ip pim sparse-mode
  ip ospf 1 area 0
!
interface Tunnel10
  description "IPSEC tunnel"
  ip address 172.16.12.1 255.255.255.0
  ip pim sparse-mode
  ip ospf 1 area 0
  tunnel source Loopback10
  tunnel mode ipsec ipv4
  tunnel destination 172.16.10.2
  tunnel protection ipsec profile ipsec_prof10
```

Loopback that is used by IPsec Tunnel in another IGP instance: the following commands establish an IPsec tunnel between the endpoints:

```
interface Loopback10
  description ipsec Loopback
  ip address 172.16.10.1 255.255.255.0
  ip ospf 2 area 0
!
```


- b) Alternatively, configure a static route for BGP neighborship, if you cannot run IGP over IPsec tunnel.

```
ip route 172.16.254.2 255.255.255.255 Tunnel10
```

Refer to the [Configuration Example for BGP EVPN VXLAN over IPsec](#)

Configuration Example for BGP EVPN VXLAN over IPsec

Figure 68: Secure VXLAN Traffic Between Two VTEPs



In this topology, VTEP1 and VTEP2 communicate through a secure VXLAN tunnel that runs through a service provider network.

The following table provides sample configurations for the devices in this topology.

VTEP1	VTEP2
-------	-------

VTEP1	VTEP2
<pre> hostname VTEP1 ! vrf definition red rd 1:100 ! address-family ipv4 route-target export 1:100 route-target import 1:100 route-target export 1:100 stitching route-target import 1:100 stitching exit-address-family ! address-family ipv6 route-target export 1:100 route-target import 1:100 route-target export 1:100 stitching route-target import 1:100 stitching exit-address-family ! ip routing ip multicast-routing ! ipv6 unicast-routing ipv6 multicast-routing ! l2vpn evpn replication-type ingress router-id Loopback1 ! l2vpn evpn instance 1500 vlan-based encapsulation vxlan replication-type ingress ! l2vpn evpn instance 1501 vlan-based encapsulation vxlan replication-type ingress ! license boot level network-advantage addon dna-advantage ! system mtu 9198 ! crypto engine compliance shield disable ! crypto ikev2 keyring ikev10_key peer mypeer address 0.0.0.0 0.0.0.0 pre-shared-key cisco123 ! crypto ikev2 profile ikev2_prof10 match identity remote address 172.16.10.2 255.255.255.255 authentication remote pre-share authentication local pre-share keyring local ikev10_key dpd 10 2 periodic ! vlan configuration 500 member vni 50000 vlan configuration 1500 member evpn-instance 1500 vni 11500 vlan configuration 1501 </pre>	<pre> hostname VTEP2 ! vrf definition red rd 1:100 ! address-family ipv4 route-target export 1:100 route-target import 1:100 route-target export 1:100 stitching route-target import 1:100 stitching exit-address-family ! address-family ipv6 route-target export 1:100 route-target import 1:100 route-target export 1:100 stitching route-target import 1:100 stitching exit-address-family ! ip routing ip multicast-routing ! ipv6 unicast-routing ipv6 multicast-routing ! l2vpn evpn replication-type ingress router-id Loopback1 ! l2vpn evpn instance 1500 vlan-based encapsulation vxlan replication-type ingress ! l2vpn evpn instance 1501 vlan-based encapsulation vxlan replication-type ingress ! license boot level network-advantage addon dna-advantage ! system mtu 9198 ! crypto engine compliance shield disable ! crypto ikev2 keyring ikev10_key peer mypeer address 0.0.0.0 0.0.0.0 pre-shared-key cisco123 ! crypto ikev2 profile ikev2_prof10 match identity remote address 172.16.10.1 255.255.255.255 authentication remote pre-share authentication local pre-share keyring local ikev10_key dpd 10 2 periodic ! vlan configuration 500 member vni 50000 vlan configuration 1500 member evpn-instance 1500 vni 11500 vlan configuration 1501 </pre>

VTEP1	VTEP2
<pre> member evpn-instance 1501 vni 11501 ! crypto ipsec transform-set tfs esp-gcm esn mode tunnel ! crypto ipsec profile ipsec_prof10 set transform-set tfs set ikev2-profile ikev2_prof10 ! ! 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 Loopback10 ip address 172.16.10.1 255.255.255.0 ip ospf 2 area 0 ! interface Tunnell0 description "IPSEC tunnel" ip address 172.16.12.1 255.255.255.0 ip pim sparse-mode ip ospf 1 area 0 tunnel source Loopback10 tunnel mode ipsec ipv4 tunnel destination 172.16.10.2 tunnel protection ipsec profile ipsec_prof10 ! ! interface TwentyFiveGigE1/0/1 ! interface TwentyFiveGigE1/0/2 ! ! interface TwentyFiveGigE1/0/3 description "Connected to VTEP2" no switchport ip address 10.3.1.1 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ! interface TwentyFiveGigE1/0/16 description "Host" switchport trunk allowed vlan 1500-1503 switchport mode trunk ! ! interface Vlan500 description "Core SVI" vrf forwarding red ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable </pre>	<pre> member evpn-instance 1501 vni 11501 ! crypto ipsec transform-set tfs esp-gcm esn mode tunnel ! crypto ipsec profile ipsec_prof10 set transform-set tfs set ikev2-profile ikev2_prof10 responder-only ! 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 Loopback10 ip address 172.16.10.2 255.255.255.255 ip ospf 2 area 0 ! interface Tunnell0 description "IPSEC tunnel" ip address 172.16.12.2 255.255.255.0 ip pim sparse-mode ip ospf 1 area 0 tunnel source Loopback10 tunnel mode ipsec ipv4 tunnel destination 172.16.10.1 tunnel protection ipsec profile ipsec_prof10 ! ! interface TwentyFiveGigE1/0/1 description "Host" switchport access vlan 1500 switchport trunk allowed vlan 1500-1503 switchport mode trunk ! interface TwentyFiveGigE1/0/3 description "connected to VTEP1" no switchport ip address 10.3.1.2 255.255.255.0 ip pim sparse-mode ip ospf network point-to-point ip ospf 2 area 0 ! ! ! ! ! interface Vlan500 description "Core SVI" vrf forwarding red ip unnumbered Loopback1 ip pim sparse-mode ipv6 enable </pre>

VTEP1	VTEP2
<pre> no autostate ! interface Vlan1500 mac-address 0000.00aa.00aa vrf forwarding red ip address 192.168.1.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ! interface Vlan1501 mac-address 0000.00bb.00bb vrf forwarding red ip address 192.168.2.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 50000 vrf red member vni 11500 ingress-replication member vni 11501 ingress-replication ! router ospf 1 ! router ospf 2 ! router bgp 1 bgp router-id interface Loopback0 bgp log-neighbor-changes bgp update-delay 1 bgp graceful-restart no bgp default ipv4-unicast neighbor 172.16.255.2 remote-as 1 neighbor 172.16.255.2 update-source Loopback0 ! address-family ipv4 redistribute static redistribute connected neighbor 172.16.255.2 activate exit-address-family ! address-family ipv6 redistribute connected redistribute static exit-address-family ! address-family l2vpn evpn neighbor 172.16.255.2 activate neighbor 172.16.255.2 send-community both exit-address-family ! address-family ipv4 vrf red advertise l2vpn evpn redistribute static redistribute connected exit-address-family ! address-family ipv6 vrf red redistribute connected </pre>	<pre> no autostate ! interface Vlan1500 mac-address 0000.00aa.00aa vrf forwarding red ip address 192.168.1.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ! interface Vlan1501 mac-address 0000.00bb.00bb vrf forwarding red ip address 192.168.2.1 255.255.255.0 ip pim sparse-mode ip igmp version 3 ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp member vni 50000 vrf red member vni 11500 ingress-replication member vni 11501 ingress-replication ! router ospf 1 ! router ospf 2 ! router bgp 1 bgp router-id interface Loopback0 bgp log-neighbor-changes bgp update-delay 1 bgp graceful-restart no bgp default ipv4-unicast neighbor 172.16.255.1 remote-as 1 neighbor 172.16.255.1 update-source Loopback0 ! address-family ipv4 redistribute static redistribute connected neighbor 172.16.255.1 activate exit-address-family ! address-family ipv6 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 red advertise l2vpn evpn redistribute static redistribute connected exit-address-family ! address-family ipv6 vrf red redistribute connected redistribute static </pre>

VTEP1	VTEP2
<pre> redistribute static advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.2 !</pre>	<pre> advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.2 !</pre>

Verifying BGP EVPN VXLAN over IPsec

The following sections provide sample output of **show** commands to verify the BGP EVPN VXLAN over IPsec configuration on the devices in the [Configuration Example for BGP EVPN VXLAN over IPsec](#).

Outputs to Verify the Configuration on VTEP1

VTEP1# **show nve peers**

```
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag
```

Interface	VNI	Type	Peer-IP	RMAC/Num_RTs	eVNI	state	flags	UP time
nve1	50000	L3CP	172.16.254.2	34ed.1b7e.44d0	50000	UP	A/M/4	00:18:51
nve1	11500	L2CP	172.16.254.2	3	11500	UP	N/A	00:18:51
nve1	11501	L2CP	172.16.254.2	3	11501	UP	N/A	00:18:51

VTEP1# **show l2vpn evpn evi detail**

```

EVPN instance:      1500 (VLAN Based)
RD:                 172.16.254.1:1500 (auto)
Import-RTs:         1:1500
Export-RTs:         1:1500
Per-EVI Label:     none
State:              Established
Replication Type:   Ingress
Encapsulation:     vxlan
IP Local Learn:    Enabled (global)
Adv. Def. Gateway: Disabled (global)
Re-originate RT5:  Disabled
Adv. Multicast:    Disabled (global)
Vlan:              1500
  Protected:       False
  Ethernet-Tag:    0
  State:           Established
  Flood Suppress:  Attached
  Core If:         Vlan500
  Access If:      Vlan1500
  NVE If:         nve1
  RMAC:           34ed.1b7e.4350
  Core Vlan:      500
  L2 VNI:         11500
  L3 VNI:         50000
  VTEP IP:        172.16.254.1
  VRF:           red
  IPv4 IRB:      Enabled
  IPv6 IRB:      Disabled
Pseudoports:
  TwentyFiveGigE1/0/16 service instance 1500
    Routes: 1 MAC, 1 MAC/IP
Peers:
  172.16.254.2
```

```

Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD

EVPN instance:      1501 (VLAN Based)
RD:                 172.16.254.1:1501 (auto)
Import-RTs:        1:1501
Export-RTs:        1:1501
Per-EVI Label:     none
State:              Established
Replication Type:   Ingress
Encapsulation:     vxlan
IP Local Learn:    Enabled (global)
Adv. Def. Gateway: Disabled (global)
Re-originate RT5:  Disabled
Adv. Multicast:    Disabled (global)
Vlan:              1501
  Protected:       False
  Ethernet-Tag:    0
  State:           Established
  Flood Suppress:  Attached
  Core IF:         Vlan500
  Access If:       Vlan1501
  NVE If:          nve1
  RMAC:           34ed.1b7e.4350
  Core Vlan:       500
  L2 VNI:          11501
  L3 VNI:          50000
  VTEP IP:         172.16.254.1
  VRF:             red
  IPv4 IRB:        Enabled
  IPv6 IRB:        Disabled
Pseudoports:
  TwentyFiveGigE1/0/16 service instance 1501
    Routes: 1 MAC, 1 MAC/IP
Peers:
  172.16.254.2
    Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD

```

VTEP1# **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

```

10.0.0.0/8 is variably subnetted, 2 subnets, 3 masks
C       10.3.1.0/24 is directly connected, TwentyFiveGigE1/0/3
L       10.3.1.1/32 is directly connected, TwentyFiveGigE1/0/3
172.16.0.0/16 is variably subnetted, 9 subnets, 2 masks
C       172.16.10.0/24 is directly connected, Loopback10
L       172.16.10.1/32 is directly connected, Loopback10
O       172.16.10.2/32 [110/2] via 10.3.1.2, 00:35:52, TwentyFiveGigE1/0/3
C       172.16.12.0/24 is directly connected, Tunnel10
L       172.16.12.1/32 is directly connected, Tunnel10
C       172.16.254.1/32 is directly connected, Loopback1

```

```
O      172.16.254.2/32 [110/1001] via 172.16.12.2, 00:29:07, Tunnel10
C      172.16.255.1/32 is directly connected, Loopback0
O      172.16.255.2/32 [110/1001] via 172.16.12.2, 00:29:07, Tunnel10
```

VTEP1# **show ip route vrf red**

Routing Table: red

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

```
          192.168.1.0/24 is variably subnetted, 4 subnets, 2 masks
C      192.168.1.0/24 is directly connected, Vlan1500
L      192.168.1.1/32 is directly connected, Vlan1500
S      192.168.1.100/32 is directly connected, Vlan1500
B      192.168.1.200/32 [200/0] via 172.16.254.2, 00:33:05, Vlan500
          192.168.2.0/24 is variably subnetted, 4 subnets, 2 masks
C      192.168.2.0/24 is directly connected, Vlan1501
L      192.168.2.1/32 is directly connected, Vlan1501
S      192.168.2.100/32 is directly connected, Vlan1501
B      192.168.2.200/32 [200/0] via 172.16.254.2, 00:01:39, Vlan500
```

VTEP1# **show bgp l2vpn evpn all**

BGP table version is 249, 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: 172.16.254.1:1500
*>i [2] [172.16.254.1:1500] [0] [48] [001201000001] [0] [*] /20
      172.16.254.2          0          100          0 ?
*>i [2] [172.16.254.1:1500] [0] [48] [001201000001] [32] [192.168.1.200] /24
      172.16.254.2          0          100          0 ?
*> [2] [172.16.254.1:1500] [0] [48] [001501000001] [0] [*] /20
      0.0.0.0          32768 ?
*> [2] [172.16.254.1:1500] [0] [48] [001501000001] [32] [192.168.1.100] /24
      0.0.0.0          32768 ?
Route Distinguisher: 172.16.254.1:1501
*>i [2] [172.16.254.1:1501] [0] [48] [001201000002] [0] [*] /20
      172.16.254.2          0          100          0 ?
*>i [2] [172.16.254.1:1501] [0] [48] [001201000002] [32] [192.168.2.200] /24
      Network          Next Hop          Metric LocPrf Weight Path
      172.16.254.2          0          100          0 ?
*> [2] [172.16.254.1:1501] [0] [48] [001501000002] [0] [*] /20
      0.0.0.0          32768 ?
*> [2] [172.16.254.1:1501] [0] [48] [001501000002] [32] [192.168.2.100] /24
      0.0.0.0          32768 ?
Route Distinguisher: 172.16.254.2:1500
```

```

*>i [2][172.16.254.2:1500][0][48][001201000001][0][*]/20
      172.16.254.2          0    100    0 ?
*>i [2][172.16.254.2:1500][0][48][001201000001][32][192.168.1.200]/24
      172.16.254.2          0    100    0 ?
Route Distinguisher: 172.16.254.2:1501
*>i [2][172.16.254.2:1501][0][48][001201000002][0][*]/20
      172.16.254.2          0    100    0 ?
*>i [2][172.16.254.2:1501][0][48][001201000002][32][192.168.2.200]/24
      172.16.254.2          0    100    0 ?
Route Distinguisher: 172.16.254.1:1500
*> [3][172.16.254.1:1500][0][32][172.16.254.1]/17
      0.0.0.0                32768 ?
*>i [3][172.16.254.1:1500][0][32][172.16.254.2]/17
      172.16.254.2          0    100    0 ?
Route Distinguisher: 172.16.254.1:1501
*> [3][172.16.254.1:1501][0][32][172.16.254.1]/17
      Network          Next Hop          Metric LocPrf Weight Path
      0.0.0.0                32768 ?
*>i [3][172.16.254.1:1501][0][32][172.16.254.2]/17
      172.16.254.2          0    100    0 ?
Route Distinguisher: 172.16.254.2:1500
*>i [3][172.16.254.2:1500][0][32][172.16.254.2]/17
      172.16.254.2          0    100    0 ?
Route Distinguisher: 172.16.254.2:1501
*>i [3][172.16.254.2:1501][0][32][172.16.254.2]/17
      172.16.254.2          0    100    0 ?
Route Distinguisher: 1:100 (default for vrf red)
* i [5][1:100][0][24][192.168.1.0]/17
      172.16.254.2          0    100    0 ?
*> [5][1:100][0][24][192.168.1.0]/17
      0.0.0.0                0    32768 ?
* i [5][1:100][0][24][192.168.2.0]/17
      172.16.254.2          0    100    0 ?
*> [5][1:100][0][24][192.168.2.0]/17
      0.0.0.0                0    32768 ?

```

VTEP1# **show crypto session detail**

Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
 K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
 X - IKE Extended Authentication, F - IKE Fragmentation
 R - IKE Auto Reconnect, U - IKE Dynamic Route Update
 S - SIP VPN

```

Interface: Tunnel10
Profile: ikev2_prof10
Uptime: 00:16:58
Session status: UP-ACTIVE
Peer: 172.16.10.2 port 500 fvrf: (none) ivrf: (none)
      Phase1_id: 172.16.10.2
      Desc: (none)
      Session ID: 3
      IKEv2 SA: local 172.16.10.1/500 remote 172.16.10.2/500 Active
      Capabilities:DU connid:1 lifetime:23:43:02
      IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0
      Active SAs: 2, origin: crypto map
      Inbound: #pkts dec'ed 1016508 drop 0 life (KB/Sec) 1058011/2581
      Outbound: #pkts enc'ed 239 drop 0 life (KB/Sec) 36/2581

```

VTEP1# **show int tunnel10 stats**

```

Tunnel10
      Switching path      Pkts In   Chars In   Pkts Out   Chars Out
      Processor           0         0           2           64
      Route cache         0         0           0           0
      Distributed cache   1056533  1092057464 484         56333

```

```
Total      1056533 1092057464      486      56397
```

Outputs to Verify the Configuration on VTEP2

VTEP2# show nve peers

```
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag
```

Interface	VNI	Type	Peer-IP	RMAC/Num RTs	eVNI	state	flags	UP time
nve1	50000	L3CP	172.16.254.1	34ed.1b7e.4350	50000	UP	A/M/4	00:20:04
nve1	11500	L2CP	172.16.254.1	3	11500	UP	N/A	00:20:04
nve1	11501	L2CP	172.16.254.1	3	11501	UP	N/A	00:20:04

VTEP2# show l2vpn evpn evi detail

```
EVPN instance:      1500 (VLAN Based)
RD:                 172.16.254.2:1500 (auto)
Import-RTs:        1:1500
Export-RTs:        1:1500
Per-EVI Label:     none
State:              Established
Replication Type:  Ingress
Encapsulation:     vxlan
IP Local Learn:    Enabled (global)
Adv. Def. Gateway: Disabled (global)
Re-originate RT5: Disabled
Adv. Multicast:    Disabled (global)
Vlan:              1500
  Protected:       False
  Ethernet-Tag:    0
  State:           Established
  Flood Suppress: Attached
  Core If:         Vlan500
  Access If:       Vlan1500
  NVE If:          nve1
  RMAC:            34ed.1b7e.44d0
  Core Vlan:       500
  L2 VNI:          11500
  L3 VNI:          50000
  VTEP IP:         172.16.254.2
  VRF:             red
  IPv4 IRB:        Enabled
  IPv6 IRB:        Disabled
Pseudoports:
  TwentyFiveGigE1/0/1 service instance 1500
    Routes: 1 MAC, 1 MAC/IP
Peers:
  172.16.254.1
    Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD

EVPN instance:      1501 (VLAN Based)
RD:                 172.16.254.2:1501 (auto)
Import-RTs:        1:1501
Export-RTs:        1:1501
Per-EVI Label:     none
State:              Established
Replication Type:  Ingress
Encapsulation:     vxlan
IP Local Learn:    Enabled (global)
Adv. Def. Gateway: Disabled (global)
Re-originate RT5: Disabled
Adv. Multicast:    Disabled (global)
```

```

Vlan:                1501
  Protected:         False
  Ethernet-Tag:      0
  State:             Established
  Flood Suppress:    Attached
  Core If:           Vlan500
  Access If:         Vlan1501
  NVE If:            nve1
  RMAC:              34ed.1b7e.44d0
  Core Vlan:         500
  L2 VNI:            11501
  L3 VNI:            50000
  VTEP IP:           172.16.254.2
  VRF:               red
  IPv4 IRB:          Enabled
  IPv6 IRB:          Disabled
  Pseudoports:
    TwentyFiveGigE1/0/1 service instance 1501
    Routes: 1 MAC, 1 MAC/IP
  Peers:
    172.16.254.1
    Routes: 1 MAC, 1 MAC/IP, 1 IMET, 0 EAD

```

VTEP2# **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

```

10.0.0.0/8 is variably subnetted, 2 subnets, 3 masks
C    10.3.1.0/24 is directly connected, TwentyFiveGigE1/0/3
L    10.3.1.2/32 is directly connected, TwentyFiveGigE1/0/3
172.16.0.0/16 is variably subnetted, 9 subnets, 2 masks
B    172.16.10.0/24 [200/0] via 172.16.255.1, 00:30:42
O    172.16.10.1/32 [110/2] via 10.3.1.1, 00:33:32, TwentyFiveGigE1/0/3
C    172.16.10.2/32 is directly connected, Loopback10
C    172.16.12.0/24 is directly connected, Tunnel10
L    172.16.12.2/32 is directly connected, Tunnel10
O    172.16.254.1/32 [110/1001] via 172.16.12.1, 00:26:48, Tunnel10
C    172.16.254.2/32 is directly connected, Loopback1
O    172.16.255.1/32 [110/1001] via 172.16.12.1, 00:26:48, Tunnel10
C    172.16.255.2/32 is directly connected, Loopback0

```

VTEP2# **show ip route vrf red**

Routing Table: red

```

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

```

192.168.1.0/24 is variably subnetted, 4 subnets, 2 masks
C    192.168.1.0/24 is directly connected, Vlan1500
L    192.168.1.1/32 is directly connected, Vlan1500
B    192.168.1.100/32 [200/0] via 172.16.254.1, 00:00:41, Vlan500
S    192.168.1.200/32 is directly connected, Vlan1500
192.168.2.0/24 is variably subnetted, 4 subnets, 2 masks
C    192.168.2.0/24 is directly connected, Vlan1501
L    192.168.2.1/32 is directly connected, Vlan1501
B    192.168.2.100/32 [200/0] via 172.16.254.1, 00:00:35, Vlan500
S    192.168.2.200/32 is directly connected, Vlan1501

```

VTEP2# **show bgp l2vpn evpn all**

BGP table version is 309, 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: 172.16.254.1:1500
*>i [2] [172.16.254.1:1500] [0] [48] [001501000001] [0] [*]/20
      172.16.254.1          0          100          0 ?
*>i [2] [172.16.254.1:1500] [0] [48] [001501000001] [32] [192.168.1.100]/24
      172.16.254.1          0          100          0 ?
Route Distinguisher: 172.16.254.1:1501
*>i [2] [172.16.254.1:1501] [0] [48] [001501000002] [0] [*]/20
      172.16.254.1          0          100          0 ?
*>i [2] [172.16.254.1:1501] [0] [48] [001501000002] [32] [192.168.2.100]/24
      172.16.254.1          0          100          0 ?
Route Distinguisher: 172.16.254.2:1500
*> [2] [172.16.254.2:1500] [0] [48] [001201000001] [0] [*]/20
      0.0.0.0                32768 ?
      Network          Next Hop          Metric LocPrf Weight Path
*> [2] [172.16.254.2:1500] [0] [48] [001201000001] [32] [192.168.1.200]/24
      0.0.0.0                32768 ?
*>i [2] [172.16.254.2:1500] [0] [48] [001501000001] [0] [*]/20
      172.16.254.1          0          100          0 ?
*>i [2] [172.16.254.2:1500] [0] [48] [001501000001] [32] [192.168.1.100]/24
      172.16.254.1          0          100          0 ?
Route Distinguisher: 172.16.254.2:1501
*> [2] [172.16.254.2:1501] [0] [48] [001201000002] [0] [*]/20
      0.0.0.0                32768 ?
*> [2] [172.16.254.2:1501] [0] [48] [001201000002] [32] [192.168.2.200]/24
      0.0.0.0                32768 ?
*>i [2] [172.16.254.2:1501] [0] [48] [001501000002] [0] [*]/20
      172.16.254.1          0          100          0 ?
*>i [2] [172.16.254.2:1501] [0] [48] [001501000002] [32] [192.168.2.100]/24
      172.16.254.1          0          100          0 ?
Route Distinguisher: 172.16.254.1:1500
*>i [3] [172.16.254.1:1500] [0] [32] [172.16.254.1]/17
      172.16.254.1          0          100          0 ?
Route Distinguisher: 172.16.254.1:1501
*>i [3] [172.16.254.1:1501] [0] [32] [172.16.254.1]/17
      172.16.254.1          0          100          0 ?

```

```

Route Distinguisher: 172.16.254.2:1500
  Network          Next Hop          Metric LocPrf Weight Path
*>i [3][172.16.254.2:1500][0][32][172.16.254.1]/17
    172.16.254.1          0      100      0 ?
*> [3][172.16.254.2:1500][0][32][172.16.254.2]/17
    0.0.0.0                32768 ?
Route Distinguisher: 172.16.254.2:1501
*>i [3][172.16.254.2:1501][0][32][172.16.254.1]/17
    172.16.254.1          0      100      0 ?
*> [3][172.16.254.2:1501][0][32][172.16.254.2]/17
    0.0.0.0                32768 ?
Route Distinguisher: 1:100 (default for vrf red)
* i [5][1:100][0][24][192.168.1.0]/17
    172.16.254.1          0      100      0 ?
*> [5][1:100][0][24][192.168.1.0]/17
    0.0.0.0                32768 ?
* i [5][1:100][0][24][192.168.2.0]/17
    172.16.254.1          0      100      0 ?
*> [5][1:100][0][24][192.168.2.0]/17
    0.0.0.0                32768 ?

```

VTEP2# **show crypto session detail**

Crypto session current status

Code: C - IKE Configuration mode, D - Dead Peer Detection
 K - Keepalives, N - NAT-traversal, T - cTCP encapsulation
 X - IKE Extended Authentication, F - IKE Fragmentation
 R - IKE Auto Reconnect, U - IKE Dynamic Route Update
 S - SIP VPN

```

Interface: Tunnel10
Profile: ikev2_prof10
Uptime: 00:17:28
Session status: UP-ACTIVE
Peer: 172.16.10.1 port 500 fvrf: (none) ivrf: (none)
  Phase1_id: 172.16.10.1
  Desc: (none)
  Session ID: 4
  IKEv2 SA: local 172.16.10.2/500 remote 172.16.10.1/500 Active
    Capabilities:DU connid:1 lifetime:23:42:32
  IPSEC FLOW: permit ip 0.0.0.0/0.0.0.0 0.0.0.0/0.0.0.0
    Active SAs: 2, origin: crypto map
    Inbound: #pkts dec'ed 245 drop 0 life (KB/Sec) 30/2552
    Outbound: #pkts enc'ed 1043067 drop 0 life (KB/Sec) 1118249/2552

```

VTEP2# **show int tunnel10 stats**

```

Tunnel10
  Switching path  Pkts In  Chars In  Pkts Out  Chars Out
  Processor       0         0         0         0
  Route cache     0         0         0         0
  Distributed cache 228      21855    1027163  1082443955
  Total           228      21855    1027163  1082443955

```




CHAPTER 15

Configuring BGP EVPN VXLANv6

This document describes new deployment and migration scenarios for implementing IPv6 in the underlay of a BGP EVPN VXLAN fabric.

- [Restrictions for BGP EVPN VXLANv6, on page 731](#)
- [Information About BGP EVPN VXLANv6, on page 731](#)
- [How to Configure BGP EVPN VXLANv6, on page 733](#)
- [Migration from VXLANv4 to VXLANv6, on page 736](#)
- [Troubleshooting BGP EVPN VXLANv6, on page 740](#)
- [Configuration Examples for BGP EVPN VXLANv6, on page 740](#)

Restrictions for BGP EVPN VXLANv6

- Optimized Layer 2 overlay multicast is not supported in BGP EVPN VXLANv6.
- BGP EVPN VXLANv6 does not support centralized anycast gateway deployment. Only distributed anycast gateway is supported.
- BGP EVPN VXLANv6 does not support the following features:
 - Multi-Homing
 - Access Virtual Forwarding Interface (VFI)
 - Fabric Netflow

Information About BGP EVPN VXLANv6

The ever-expanding needs of the enterprise networks to enhance the IP addressing capabilities and simplify the network configurations makes it pertinent to adopt IPv6-enabled networks. The larger IPv6 address space allows networks to scale and provide global reachability. The simplified IPv6 packet header format handles packets more efficiently. IPv6 prefix aggregation, simplified network renumbering, and IPv6 site multihoming capabilities provide an IPv6 addressing hierarchy that allows for more efficient routing. When you enable IPv6 in the underlay for VXLAN tunnel endpoints (VTEPS), the overlay traffic is encapsulated with the IPv6 VXLAN encapsulation. IPv6 connectivity between the VTEPS is established by enabling IPv6 unicast routing protocols. For overlay multi-destination traffic (BUM), when underlay multicast is used, IPv6 multicast protocol must be enabled.

For a new deployment, you can build your BGP EVPN VXLAN fabric with IPv6 underlay. For an existing BGP EVPN VXLAN fabric with IPv4 underlay, you can seamlessly migrate to an IPv6 only or dual stack underlay.

BGP EVPN VXLANv6 Fabric Support

In Cisco IOS XE Dublin 17.10.1 and later releases, a BGP EVPN VXLAN fabric with an IPv6 underlay supports the following features:

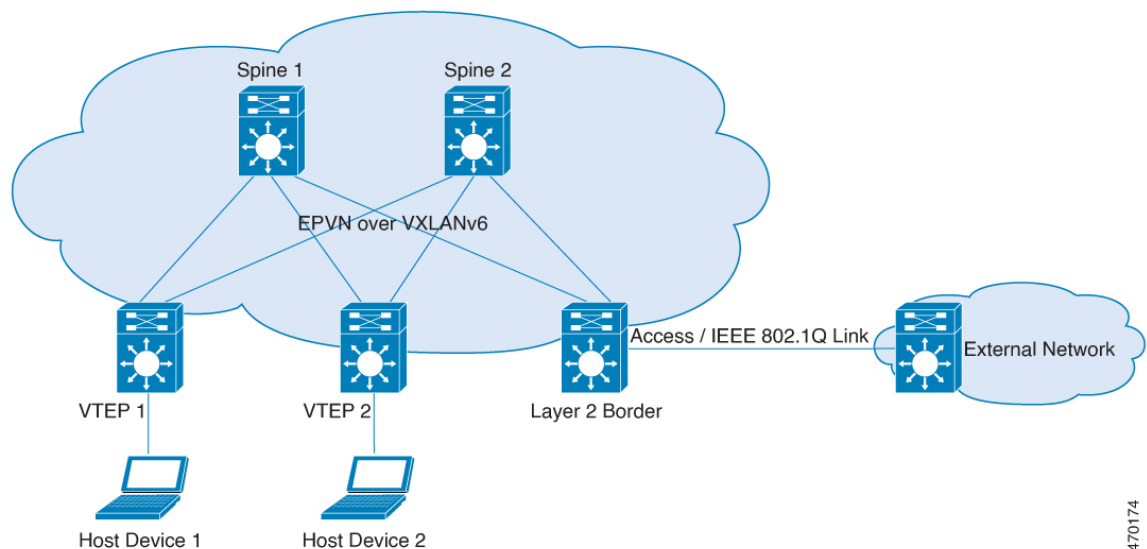
- Ingress Replication or Static Multicast Replication
- Layer 2 Gateway
- Distributed Anycast Gateway with Symmetric Integrated Routing and Bridging
- Layer 2 External connectivity with IEEE 802.1Q network
- Layer 3 External Connectivity with VRF-Lite or MPLS Layer 3 VPN network
- Seamless migration between EVPN VXLANv4 fabric and EVPN VXLANv6 fabric
- IPv6 BGP EVPN Neighbor Peering

EVPN VXLANv6 Layer 2 Overlay

The Layer 2 Gateway with VXLANv6 deployment supports IPv6 transport in the underlay. The VXLAN tunnels and BGP peering between Layer 2 Gateways are both IPv6 based. The overlay has only Layer 2 routes.

A Layer 2 Gateway with VXLANv6 supports Layer 2 bridging between Layer 2 Gateways. Layer 3 Routing is implemented with the external Layer 3 Gateways.

Figure 69: EVPN VXLANv6 Fabric with Layer 2 Gateway



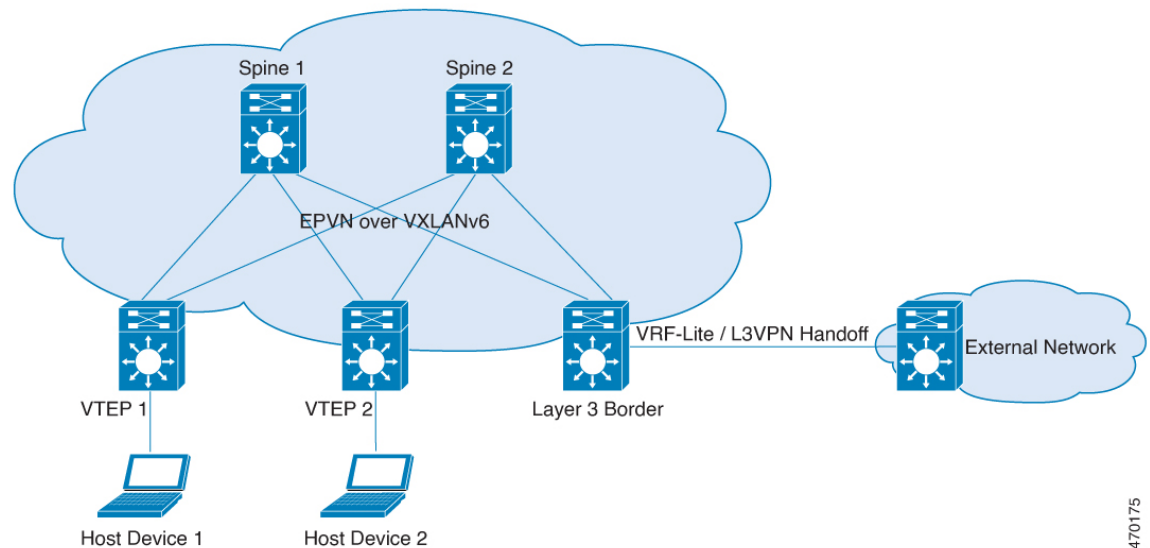
The topology in the [Figure 69: EVPN VXLANv6 Fabric with Layer 2 Gateway](#) shows a BGP EVPN VXLAN fabric with a Layer 2 overlay that supports IPv6 underlay. The fabric supports both IPv4 and IPv6 hosts.

EVPN VXLANv6 Overlay with Distributed Anycast Gateway

In a BGP EVPN VXLANv6 fabric with Distributed Anycast Gateway, the underlay supports IPv6 transport. The VXLAN tunnels and BGP peering, both support IPv6 addressing. The overlay has both Layer 2 and Layer 3 routes and supports MAC, IPv4 and IPv6 addressing.

Distributed Anycast Gateway with VXLANv6 supports Layer 2 bridging, Layer 3 routing within the EVPN fabric, and Layer 3 stitching through Layer 3 border gateway with IPv6 underlay. For Layer 3 routing, both IPv4 and IPv6 overlay prefixes are supported over IPv6 underlay.

Figure 70: BGP EVPN VXLANv6 Fabric with Distributed Anycast Gateway



Topology in the [Figure 70: BGP EVPN VXLANv6 Fabric with Distributed Anycast Gateway](#) shows a BGP EVPN VXLANv6 fabric with integrated routing and bridging (IRB) using distributed anycast gateway (DAG). The fabric supports both IPv4 and IPv6 hosts, and IPv6 BGP peering.

How to Configure BGP EVPN VXLANv6

This section describes only the changes in the underlay configuration. For the overlay and complete configuration of the fabric, refer [How to Configure EVPN VXLAN Layer 2 Overlay Network](#), on page 16 and [Information About EVPN VXLAN Integrated Routing and Bridging](#), on page 95.

Configure VXLANv6 for BGP EVPN Fabric with Layer 2 Overlay

Implement the following tasks to configure IPv6 underlay for a BGP EVPN VXLAN fabric with Layer 2 overlay.

Procedure

- Step 1** Configure an explicit router-id for L2VPN EVPN.

```

l2vpn
 router-id 1.1.1.3
l2vpn evpn
 replication-type ingress
end

```

Step 2 Configure an L2EVPN instance.

```

l2vpn evpn instance 1 vlan-based
encapsulation vxlan

```

Step 3 Configure Loopback interface.

```

interface Loopback1
 ipv6 address ABCD:1::2/128
 ipv6 ospf 1 area 0

```

Step 4 Configure NVE interface.

```

interface nve1
 source-interface Loopback1
 host-reachability protocol bgp
 vxlan encapsulation ipv6
 member vni 20011 ingress-replication
 member vni 20012 mcast-group FF0E::12

```

Step 5 Configure the VLAN.

```

vlan configuration 11
 member evpn-instance 1 vni 20011

```

Step 6 Configure the L2VPN EVPN BGP peering.

```

interface Loopback0
 description BGP UPDATE SOURCE
 ipv6 address ABCD:1::1/128
 ipv6 ospf 1 area 0
!
router bgp 100
 bgp router-id 1.1.1.1
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor ABCD:99::99 remote-as 100
 neighbor ABCD:99::99 update-source Loopback0
!
 address-family l2vpn evpn
  neighbor ABCD:99::99 activate
  neighbor ABCD:99::99 send-community both
 exit-address-family
!

```

Configure VXLANv6 for BGP EVPN Fabric with Distributed Anycast Gateway

Implement the following tasks to configure IPv6 underlay for a BGP EVPN VXLAN fabric with a Distributed Anycast Gateway (DAG).

Procedure

Step 1 Configure IP-VRF.

```
vrf definition red
 rd 100:1
 !
 address-family ipv4
  route-target export 100:100
  route-target import 100:100
  route-target export 100:100 stitching
  route-target import 100:100 stitching
 exit-address-family
 !
 address-family ipv6
 route-target export 100:200
  route-target import 100:200
 route-target export 100:200 stitching
  route-target import 100:200 stitching
 exit-address-family
end
```

Step 2 Configure Access SVI IRB interface.

```
interface Vlan11
 mac-address 0011.0011.0011
 vrf forwarding red
 ip address 192.168.11.254 255.255.255.0
 ipv6 address 2001:11::254/64
end
```

Step 3 Configure IPv6 BGP peering.

```
router bgp 100
 bgp router-id 1.1.1.1
 bgp log-neighbor-changes
 bgp graceful-restart
 neighbor ABCD:99::99 remote-as 100
 neighbor ABCD:99::99 update-source Loopback0
 !
 address-family l2vpn evpn
  neighbor ABCD:99::99 activate
  neighbor ABCD:99::99 send-community both
 exit-address-family
 !
 address-family ipv4 vrf red
  advertise l2vpn evpn
  redistribute connected
 exit-address-family
 address-family ipv6 vrf red
  advertise l2vpn evpn
  redistribute connected
 exit-address-family
end
```

Step 4 Configure NVE interface.

```
interface nve1
 source-interface Loopback1
 host-reachability protocol bgp
 vxlan encapsulation ipv6
 member vni 30001 vrf red
```

```
member vni 20011 ingress-replication
member vni 20012 mcast-group ff0e::12
```

Step 5 Configure the core SVI IRB interface.

```
vlan configuration 3
member vni 30001
interface Vlan3
vrf forwarding red
ip unnumbered Loopback192
ip pim sparse-mode
ipv6 unnumbered Loopback192
no autostate
end
```

Migration from VXLANv4 to VXLANv6

You can seamlessly migrate your existing network from VXLANv4 to VXLANv6 underlay. To achieve a nondisruptive migration, the EVPN VXLAN network should migrate incrementally from IPv4 to IPv6 underlay. During migration, the part of the network that is migrated to IPv6 underlay remains connected with the part of the network that still works with IPv4 underlay.

The incremental migration is achieved by individually making each VTEP dual stack-capable. A dual stack VTEP node has two VTEP addresses (IPv4 and IPv6) associated with the same VXLAN Network Identifier (VNI). During the underlay migration, both these VTEP IP addresses are advertised to peers in a single BGP EVPN dual next-hop update. The receiving nodes in the underlay can choose either IPv4 or IPv6 for traffic forwarding. This ensures connectivity between the IPv4 VTEPs, the dual stack VTEPs and IPv6 VTEPs. After all the VTEPs in the fabric are made dual stack-capable, each VTEP is migrated to IPv6.

Step 1

- Configure IPv4/IPv6 Loopback interface.

```
VTEP(config)#interface Loopback1
VTEP(config-if)# ipv6 address
ABCD:1::2/128
VTEP(config-if)# ipv6 ospf 1
area 0
```

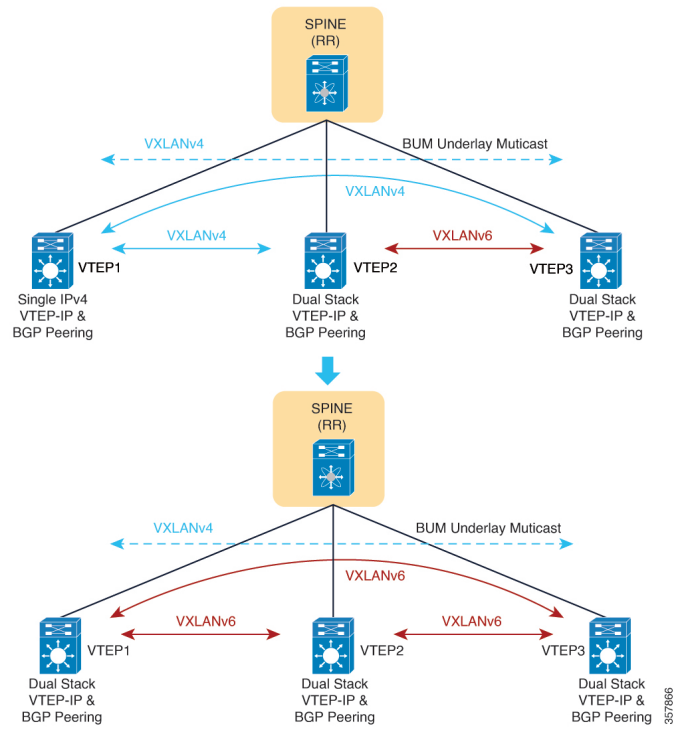
- Enable Dual Stack.

```
VTEP(config)# interface nve1
VTEP(config-if)# vxlan
encapsulation dual-stack
prefer-ipv6
```

- For unicast traffic and BUM traffic using ingress replication, the dual stack VTEP communicates with other IPv4 VTEPs using VXLANv4 and with other dual stack VTEPs using VXLANv6.

Configure static multicast replication for L2VNI.

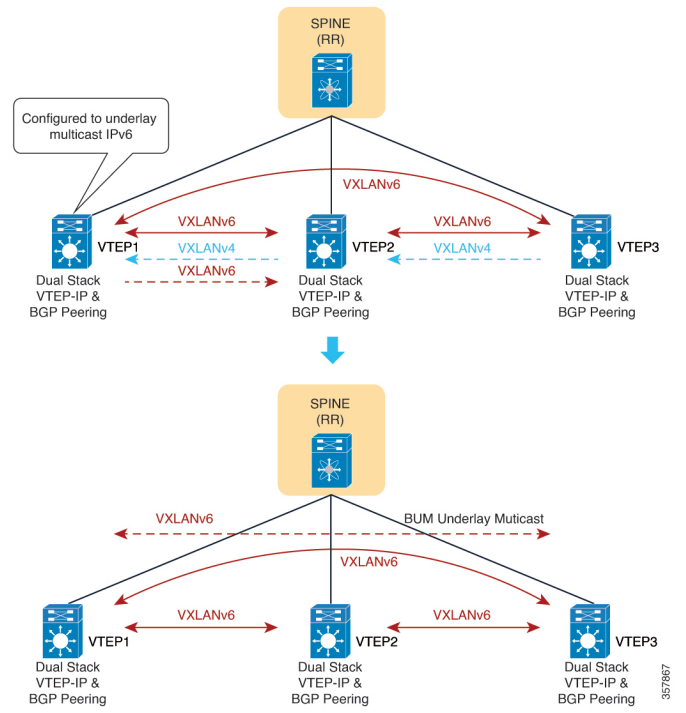
```
VTEP(config-if)# member vni
20011 mcast-group 239.2.1.11
FF0E::11
```



Step 2

- This step is only needed for BUM traffic using Multicast Replication.
- After all the VTEPs are dual stack-capable, underlay multicast switches to IPv6 Multicast.

```
VTEP(config)# interface nve1
VTEP(config-if)# vxlan
encapsulation dual-stack
prefer-ipv6 underlay-mcast ipv6
```

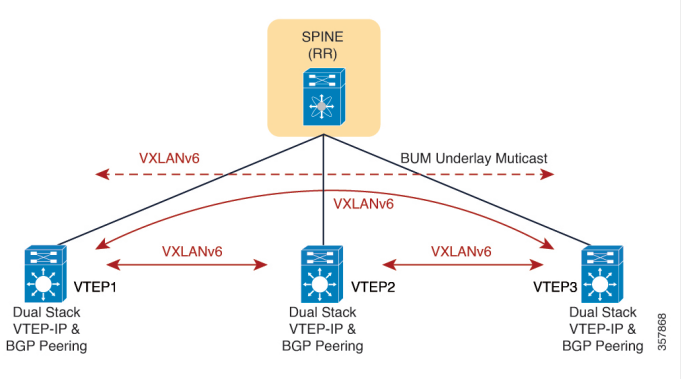
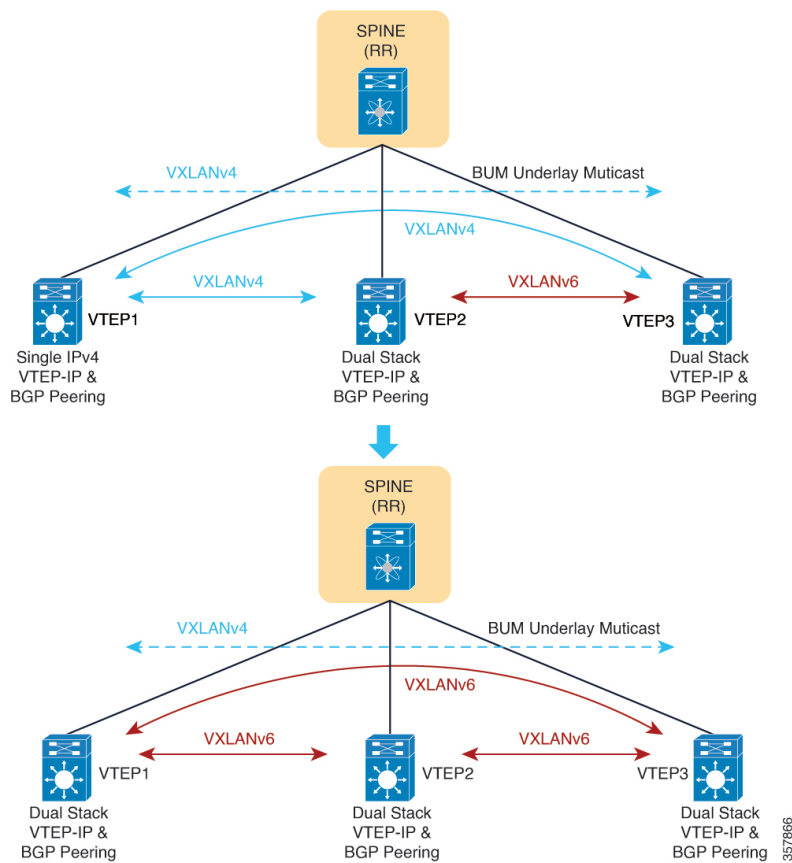


Step 3

- Configure all VTEPs to support only IPv6.

```
VTEP (config)# interface nve1
VTEP (config-if)# member vni
20011 mcast-group FF0E::11
VTEP (config-if)# vxlan
encapsulation ipv6
```

- The fabric switches to IPv6 encapsulations.

**Step 1**

1. Configure IPv4/IPv6 Loopback interface.

```
VTEP (config)# interface Loopback1
VTEP (config-if)# ipv6 address ABCD:1::2/128
VTEP (config-if)# ipv6 ospf 1 area 0
```

2. Enable Dual Stack.

```
VTEP (config)# interface nve1
VTEP (config-if)# vxlan encapsulation dual-stack prefer-ipv6
```

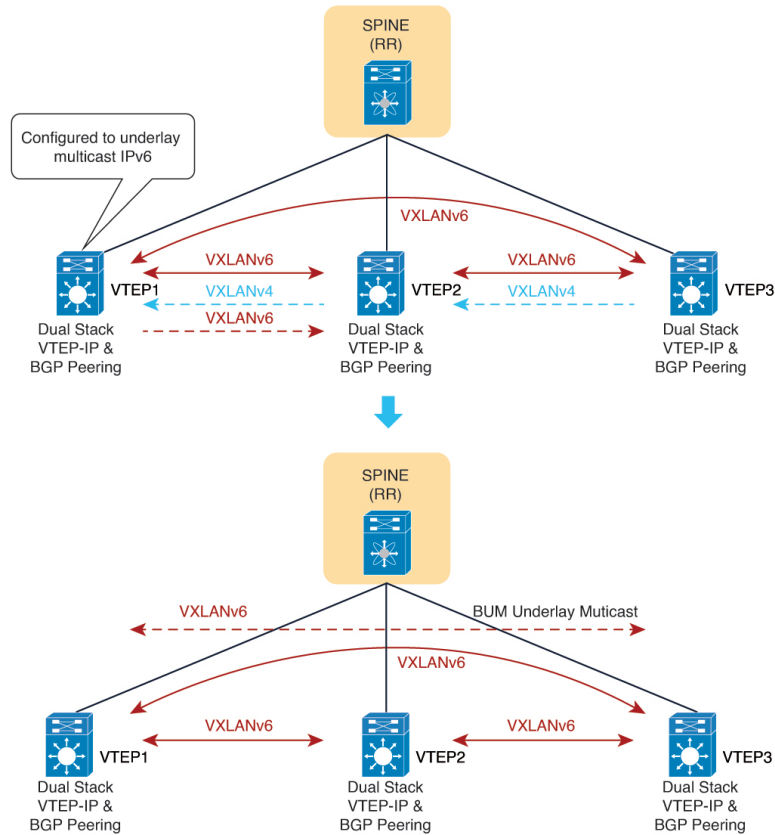

- For unicast traffic and BUM traffic using ingress replication, the dual stack VTEP communicates with other IPv4 VTEPs using VXLANv4 and with other dual stack VTEPs using VXLANv6.

Configure static multicast replication for L2VNI.

```
VTEP(config-if)# member vni 20011 mcast-group 239.2.1.11 FF0E::11
```

Step 2


This step is only needed for BUM traffic using Multicast Replication.



- After all the VTEPs are dual stack-capable, underlay multicast switches to IPv6 Multicast.

```
VTEP (config)# interface nve1
VTEP (config-if)# vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv6
```

Step 3



- Configure all VTEPs to support only IPv6.

```
VTEP(config)# interface nve1
VTEP(config-if)# member vni 20011 mcast-group FF0E::11
VTEP(config-if)# vxlan encapsulation ipv6
```

- The fabric switches to IPv6 encapsulations.

See [Migrate EVPN VXLAN to IPv6 Underlay on Catalyst 9000 Switches](#) document for detailed steps to migrate from VXLANv4 to VXLANv6.

Troubleshooting BGP EVPN VXLANv6

See [Troubleshoot EVPN VXLANv6 on Catalyst 9000 Series Switches](#) document to learn how to troubleshoot issues with the BGP EVPN VXLANv6 fabric.

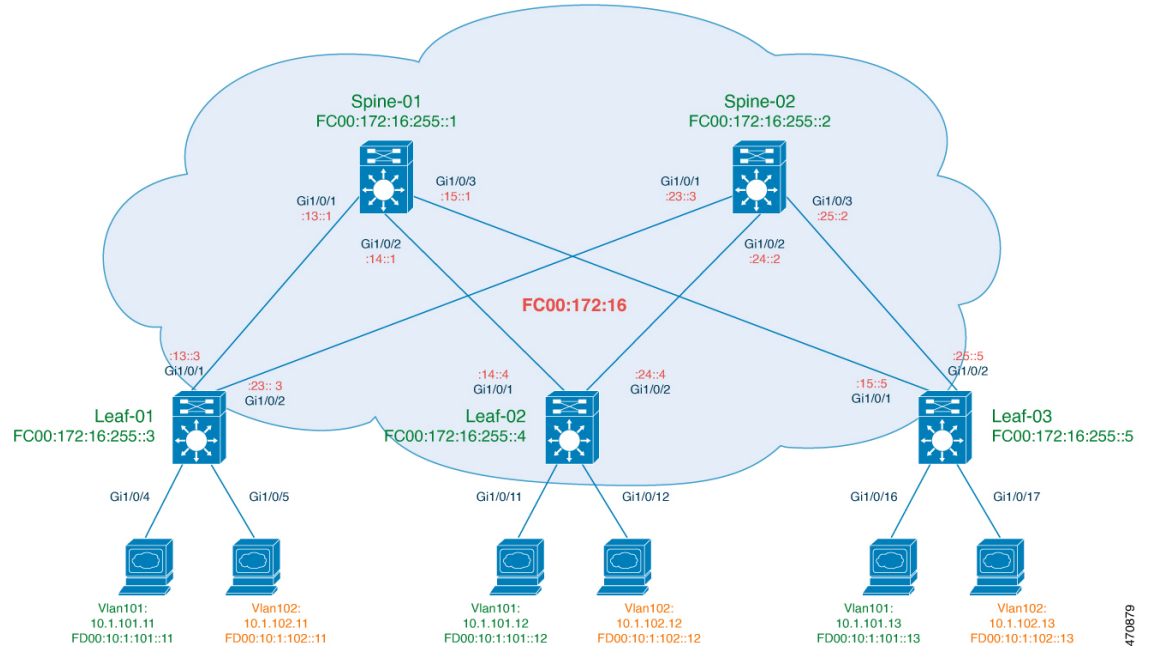
Configuration Examples for BGP EVPN VXLANv6

The following sections provide examples to show how to configure BGP EVPN VXLAN fabric with an IPv6 underlay and a dual stack underlay.

Example: Configuring BGP EVPN VXLAN with IPv6 Underlay

This example shows how to configure a BGP EVPN VXLAN fabric with only IPv6 in the underlay. The following tables provide sample configuration of the VTEPs (Leaf-01, Leaf-02, Leaf-03) and the Spine nodes (Spine-01, Spine-02) in the topology [Figure 71: BGP EVPN VXLAN with IPv6 Underlay](#).

Figure 71: BGP EVPN VXLAN with IPv6 Underlay



47/0879

Table 63: Configure the VTEPs

VTEP 1	VTEP 2	VTEP 3
--------	--------	--------

VTEP 1	VTEP 2	VTEP 3
<pre>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 ! ipv6 unicast-routing ipv6 multicast-routing ! l2vpn evpn router-id Loopback1 replication-type static default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan replication-type ingress ! 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 ipv6 address fc00:172:16:255::3/128 ipv6 ospf 1 area 0 ! interface Loopback1 ipv6 address fc00:172:16:254::3/128</pre>	<pre>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 multicast-routing ! l2vpn evpn router-id Loopback1 replication-type static default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan replication-type ingress ! 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 ipv6 address fc00:172:16:255::4/128 ipv6 ospf 1 area 0 ! interface Loopback1 ipv6 address fc00:172:16:254::4/128</pre>	<pre>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 ! ipv6 unicast-routing ipv6 multicast-routing ! l2vpn evpn router-id Loopback1 replication-type static default-gateway advertise ! l2vpn evpn instance 101 vlan-based encapsulation vxlan ! l2vpn evpn instance 102 vlan-based encapsulation vxlan replication-type ingress ! 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 ipv6 address fc00:172:16:255::5/128 ipv6 ospf 1 area 0 ! interface Loopback1 ipv6 address</pre>

Example: Configuring BGP EVPN VXLAN with IPv6 Underlay

VTEP 1	VTEP 2	VTEP 3
<pre> ip address 172.16.254.3 255.255.255.255 ipv6 ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ipv6 address fc00:172:16:13::3/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/2 no switchport ipv6 address fc00:172:16:23::3/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/4 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/5 switchport access vlan 102 switchport mode access ! ipv6 router ospf 1 router-id 172.16.255.3 ! router bgp 65001 bgp router-id 172.16.255.3 neighbor fc00:172:16:255::1 remote-as 65001 neighbor fc00:172:16:255::1 update-source Loopback0 neighbor fc00:172:16:255::2 remote-as 65001 neighbor fc00:172:16:255::2 update-source Loopback0 ! address-family l2vpn evpn neighbor fc00:172:16:255::1 activate neighbor fc00:172:16:255::1 send-community both neighbor fc00:172:16:255::2 activate neighbor fc00:172:16:255::2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected exit-address-family ! address-family ipv6 vrf green redistribute connected </pre>	<pre> ip address 172.16.254.4 255.255.255.255 ipv6 ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ipv6 address fc00:172:16:14::4/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/2 no switchport ipv6 address fc00:172:16:24::4/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/11 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/12 switchport access vlan 102 switchport mode access ! ipv6 router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp router-id 172.16.255.4 neighbor fc00:172:16:255::1 remote-as 65001 neighbor fc00:172:16:255::1 update-source Loopback0 neighbor fc00:172:16:255::2 remote-as 65001 neighbor fc00:172:16:255::2 update-source Loopback0 ! address-family l2vpn evpn neighbor fc00:172:16:255::1 activate neighbor fc00:172:16:255::1 send-community both neighbor fc00:172:16:255::2 activate neighbor fc00:172:16:255::2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected exit-address-family ! </pre>	<pre> fc00:172:16:254::5/128 ip address 172.16.254.5 255.255.255.255 ipv6 ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ipv6 address fc00:172:16:15::5/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/2 no switchport ipv6 address fc00:172:16:25::5/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/16 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/17 switchport access vlan 102 switchport mode access ! ipv6 router ospf 1 router-id 172.16.255.5 ! router bgp 65001 bgp router-id 172.16.255.5 neighbor fc00:172:16:255::1 remote-as 65001 neighbor fc00:172:16:255::1 update-source Loopback0 neighbor fc00:172:16:255::2 remote-as 65001 neighbor fc00:172:16:255::2 update-source Loopback0 ! address-family l2vpn evpn neighbor fc00:172:16:255::1 activate neighbor fc00:172:16:255::1 send-community both neighbor fc00:172:16:255::2 activate neighbor fc00:172:16:255::2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected exit-address-family </pre>

VTEP 1	VTEP 2	VTEP 3
<pre> advertise l2vpn evpn exit-address-family ! int vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address fd00:10:1:101::1/64 ipv6 enable ! int vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ipv6 address fd00:10:1:102::1/64 ipv6 enable ! int vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp vxlan encapsulation ipv6 member vni 10101 mcast-group FF0E:225::101 member vni 10102 ingress-replication member vni 50901 vrf green ! ipv6 pim rp-address fc00:172:16:255::255 ! end </pre>	<pre> address-family ipv6 vrf green ! redistribute connected advertise l2vpn evpn exit-address-family ! int vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address fd00:10:1:101::1/64 ipv6 enable ! int vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ipv6 address fd00:10:1:102::1/64 ipv6 enable ! int vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp vxlan encapsulation ipv6 member vni 10101 mcast-group FF0E:225::101 member vni 10102 ingress-replication member vni 50901 vrf green ! ipv6 pim rp-address fc00:172:16:255::255 ! end </pre>	<pre> ! address-family ipv6 vrf green redistribute connected advertise l2vpn evpn exit-address-family ! int vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address fd00:10:1:101::1/64 ipv6 enable ! int vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ipv6 address fd00:10:1:102::1/64 ipv6 enable ! int vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp vxlan encapsulation ipv6 member vni 10101 mcast-group FF0E:225::101 member vni 10102 ingress-replication member vni 50901 vrf green ! ipv6 pim rp-address fc00:172:16:255::255 ! end </pre>

Table 64: Configure the Spine Nodes

Spine Switch 1	Spine Switch 2
----------------	----------------

Spine Switch 1	Spine Switch 2
<pre> Spine-01# show running-config hostname Spine-01 ! ip routing ! ipv6 unicast-routing ipv6 multicast-routing ! system mtu 9198 ! interface Loopback0 ipv6 address fc00:172:16:255::1/128 ipv6 ospf 1 area 0 ! interface Loopback1 ipv6 address fc00:172:16:254::1/128 ipv6 ospf 1 area 0 ! interface Loopback2 ipv6 address fc00:172:16:255::255/128 ipv6 ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ipv6 address fc00:172:16:13::1/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/2 no switchport ipv6 address fc00:172:16:14::1/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/3 no switchport ipv6 address fc00:172:16:15::1/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! ipv6 router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp router-id 172.16.255.1 neighbor fc00:172:16:255::2 remote-as 65001 neighbor fc00:172:16:255::2 update-source Loopback0 neighbor fc00:172:16:255::3 remote-as 65001 neighbor fc00:172:16:255::3 update-source Loopback0 neighbor fc00:172:16:255::4 remote-as 65001 neighbor fc00:172:16:255::4 update-source Loopback0 neighbor fc00:172:16:255::5 remote-as 65001 neighbor fc00:172:16:255::5 update-source Loopback0 ! address-family l2vpn evpn </pre>	<pre> Spine-02# show running-config hostname Spine-02 ! ip routing ! ipv6 unicast-routing ipv6 multicast-routing ! system mtu 9198 ! interface Loopback0 ipv6 address fc00:172:16:255::2/128 ipv6 ospf 1 area 0 ! interface Loopback1 ipv6 address fc00:172:16:254::2/128 ipv6 ospf 1 area 0 ! interface Loopback2 ipv6 address fc00:172:16:255::255/128 ipv6 ospf 1 area 0 ! interface GigabitEthernet1/0/1 no switchport ipv6 address fc00:172:16:23::2/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/2 no switchport ipv6 address fc00:172:16:24::2/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/3 no switchport ipv6 address fc00:172:16:25::2/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! ipv6 router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp router-id 172.16.255.2 neighbor fc00:172:16:255::1 remote-as 65001 neighbor fc00:172:16:255::1 update-source Loopback0 neighbor fc00:172:16:255::3 remote-as 65001 neighbor fc00:172:16:255::3 update-source Loopback0 neighbor fc00:172:16:255::4 remote-as 65001 neighbor fc00:172:16:255::4 update-source Loopback0 neighbor fc00:172:16:255::5 remote-as 65001 neighbor fc00:172:16:255::5 update-source Loopback0 ! address-family l2vpn evpn </pre>

Spine Switch 1	Spine Switch 2
<pre>neighbor fc00:172:16:255::2 activate neighbor fc00:172:16:255::2 send-community both neighbor fc00:172:16:255::2 route-reflector-client neighbor fc00:172:16:255::3 activate neighbor fc00:172:16:255::3 send-community both neighbor fc00:172:16:255::3 route-reflector-client neighbor fc00:172:16:255::4 activate neighbor fc00:172:16:255::4 send-community both neighbor fc00:172:16:255::4 route-reflector-client neighbor fc00:172:16:255::5 activate neighbor fc00:172:16:255::5 send-community both neighbor fc00:172:16:255::5 route-reflector-client exit-address-family ! ipv6 pim rp-address fc00:172:16:255::255 ! ipv6 pim anycast-rp fc00:172:16:255::255 fc00:172:16:254::2 ! end</pre>	<pre>neighbor fc00:172:16:255::1 activate neighbor fc00:172:16:255::1 send-community both neighbor fc00:172:16:255::1 route-reflector-client neighbor fc00:172:16:255::3 activate neighbor fc00:172:16:255::3 send-community both neighbor fc00:172:16:255::3 route-reflector-client neighbor fc00:172:16:255::4 activate neighbor fc00:172:16:255::4 send-community both neighbor fc00:172:16:255::4 route-reflector-client neighbor fc00:172:16:255::5 activate neighbor fc00:172:16:255::5 send-community both neighbor fc00:172:16:255::5 route-reflector-client exit-address-family ! ipv6 pim rp-address fc00:172:16:255::255 ! ipv6 pim anycast-rp fc00:172:16:255::255 fc00:172:16:254::1 ! end</pre>

Verifying BGP EVPN VXLAN with IPv6 Underlay Configuration

The following sections provide sample output of **show** commands to verify the BGP EVPN VXLAN configuration with IPv6 underlay.

- [Outputs to Verify Configuration on VTEP 1, on page 748](#)
- [Outputs to Verify Configuration on VTEP 2, on page 754](#)
- [Outputs to Verify Configuration on VTEP 3, on page 760](#)
- [Outputs to Verify Configuration on Spine 1, on page 767](#)
- [Outputs to Verify Configuration on Spine 2, on page 772](#)

Outputs to Verify Configuration on VTEP 1

```
Leaf-01# show ipv6 route
IPv6 Routing Table - default - 20 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, la - LISP away, le - LISP extranet-policy
       lp - LISP publications, ls - LISP destinations-summary
C    FC00:172:16:13::/64 [0/0]
```

```

    via GigabitEthernet1/0/1, directly connected
L  FC00:172:16:13::3/128 [0/0]
    via GigabitEthernet1/0/1, receive
O  FC00:172:16:14::/64 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
O  FC00:172:16:15::/64 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
C  FC00:172:16:23::/64 [0/0]
    via GigabitEthernet1/0/2, directly connected
L  FC00:172:16:23::3/128 [0/0]
    via GigabitEthernet1/0/2, receive
O  FC00:172:16:24::/64 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
O  FC00:172:16:25::/64 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
O  FC00:172:16:254::1/128 [110/1]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
O  FC00:172:16:254::2/128 [110/1]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
LC FC00:172:16:254::3/128 [0/0]
    via Loopback1, receive
O  FC00:172:16:254::4/128 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
O  FC00:172:16:254::5/128 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
O  FC00:172:16:255::1/128 [110/1]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
O  FC00:172:16:255::2/128 [110/1]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
LC FC00:172:16:255::3/128 [0/0]
    via Loopback0, receive
O  FC00:172:16:255::4/128 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
O  FC00:172:16:255::5/128 [110/2]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
O  FC00:172:16:255::255/128 [110/1]
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2E4, GigabitEthernet1/0/2
L  FF00::/8 [0/0]
    via Null0, receive

```

Leaf-01# **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

```

(*, FF0E:225::101), 00:05:08/never, RP FC00:172:16:255::255, flags: SCJ
  Incoming interface: GigabitEthernet1/0/2
  RPF nbr: FE80::7E21:DFE:FE92:B2E4
  Immediate Outgoing interface list:
    Tunnell1, Forward, 00:05:08/never

```

```
(FC00:172:16:254::3, FF0E:225::101), 00:04:36/00:03:01, flags: SFJT
Incoming interface: Loopback1
RPF nbr: FE80::12B3:D5FF:FE6A:8F80
Immediate Outgoing interface list:
  GigabitEthernet1/0/1, Forward, 00:04:36/00:03:02
Inherited Outgoing interface list:
  Tunnel1, Forward, 00:05:08/never

(FC00:172:16:254::4, FF0E:225::101), 00:04:28/00:02:45, flags: SJT
Incoming interface: GigabitEthernet1/0/1
RPF nbr: FE80::7E21:DFE:FE92:B2E4
Inherited Outgoing interface list:
  Tunnel1, Forward, 00:05:08/never

(FC00:172:16:254::5, FF0E:225::101), 00:04:22/00:00:12, flags: SJT
Incoming interface: GigabitEthernet1/0/2
RPF nbr: FE80::7E21:DFE:FEBD:2CE4
Inherited Outgoing interface list:
  Tunnel1, Forward, 00:05:08/never
```

Leaf-01# **show nve peer**

```
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag
```

Interface	VNI	Type	Peer-IP	RMAC/Num_RTs	eVNI	state	flags	UP	time
nve1	50901	L3CP	FC00:172:16:254::5 \	7c21.0dbd.2748	50901	UP	A/-/4	00:03:30	
nve1	50901	L3CP	FC00:172:16:254::4 \	7c21.0dbd.9548	50901	UP	A/-/4	00:03:30	
nve1	50901	L3CP	FC00:172:16:254::5 \	7c21.0dbd.2748	50901	UP	A/M/6	00:03:30	
nve1	50901	L3CP	FC00:172:16:254::4 \	7c21.0dbd.9548	50901	UP	A/M/6	00:03:30	
nve1	10101	L2CP	FC00:172:16:254::4 \	4	10101	UP	N/A	00:03:29	
nve1	10101	L2CP	FC00:172:16:254::5 \	4	10101	UP	N/A	00:03:29	
nve1	10102	L2CP	FC00:172:16:254::4 \	6	10102	UP	N/A	00:03:29	
nve1	10102	L2CP	FC00:172:16:254::5 \	6	10102	UP	N/A	00:03:29	

Leaf-01# **show bgp l2vpn evpn summary**

```
BGP router identifier 172.16.255.3, local AS number 65001
BGP table version is 49, main routing table version 49
46 network entries using 17664 bytes of memory
62 path entries using 14384 bytes of memory
18/18 BGP path/bestpath attribute entries using 5328 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
14 BGP extended community entries using 744 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 38280 total bytes of memory
BGP activity 56/0 prefixes, 82/6 paths, scan interval 60 secs
46 networks peaked at 16:39:55 Oct 31 2022 UTC (00:01:35.087 ago)
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
FC00:172:16:255::1	4	65001	27	15	49	0	0	00:04:40	16
FC00:172:16:255::2	4	65001	27	15	49	0	0	00:04:30	16

```

Leaf-01# show bgp 12vpn evpn
BGP table version is 49, 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] [FD00:10:1:101::1]/36
      ::                               32768 ?
*>i [2] [172.16.254.3:101] [0] [48] [44D3CA286CC1] [0] [*]/20
      FC00:172:16:254::4
      0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [44D3CA286CC3] [0] [*]/20
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [7C210DBD2741] [32] [10.1.101.1]/24
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
      FC00:172:16:254::4
      0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
      FC00:172:16:254::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 ?
*> [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
      ::                               32768 ?
Route Distinguisher: 172.16.254.3:102
*> [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24
      ::                               32768 ?
*> [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
      ::                               32768 ?
*>i [2] [172.16.254.3:102] [0] [48] [44D3CA286CC2] [0] [*]/20
      FC00:172:16:254::4
      0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
      FC00:172:16:254::4
      0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [44D3CA286CC4] [0] [*]/20
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [7C210DBD954D] [32] [10.1.102.1]/24
      FC00:172:16:254::4

```

```

0 100 0 ?
*>i [2][172.16.254.3:102][0][48][7C210DBD954D][128][FD00:10:1:102::1]/36
FC00:172:16:254::4
0 100 0 ?
*> [2][172.16.254.3:102][0][48][F4CFE24334C2][0][*]/20
:: 32768 ?
*> [2][172.16.254.3:102][0][48][F4CFE24334C2][128][FE80::F6CF:E2FF:FE43:34C2]/36
:: 32768 ?
Route Distinguisher: 172.16.254.4:101
* i [2][172.16.254.4:101][0][48][44D3CA286CC1][0][*]/20
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
* i [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
* i [2][172.16.254.4:101][0][48][7C210DBD9541][128][FD00:10:1:101::1]/36
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][0][*]/20
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][32][10.1.102.12]/24
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
* i [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
* i [2][172.16.254.4:102][0][48][7C210DBD954D][128][FD00:10:1:102::1]/36
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
Route Distinguisher: 172.16.254.5:101
* i [2][172.16.254.5:101][0][48][44D3CA286CC3][0][*]/20
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
* i [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
* i [2][172.16.254.5:101][0][48][7C210DBD2741][128][FD00:10:1:101::1]/36
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [2][172.16.254.5:102][0][48][44D3CA286CC4][0][*]/20
FC00:172:16:254::5

```

```

* > i          FC00:172:16:254::5          0    100    0 ?
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
      FC00:172:16:254::5          0    100    0 ?
* > i          FC00:172:16:254::5          0    100    0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
      FC00:172:16:254::5          0    100    0 ?
* > i          FC00:172:16:254::5          0    100    0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
      FC00:172:16:254::5          0    100    0 ?
* > i          FC00:172:16:254::5          0    100    0 ?
Route Distinguisher: 172.16.254.3:102
* > [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
      ::                          32768 ?
* > i [3] [172.16.254.3:102] [0] [32] [172.16.254.4]/17
      FC00:172:16:254::4          0    100    0 ?
* > i [3] [172.16.254.3:102] [0] [32] [172.16.254.5]/17
      FC00:172:16:254::5          0    100    0 ?
Route Distinguisher: 172.16.254.4:102
* i [3] [172.16.254.4:102] [0] [32] [172.16.254.4]/17
      FC00:172:16:254::4          0    100    0 ?
* > i          FC00:172:16:254::4          0    100    0 ?
Route Distinguisher: 172.16.254.5:102
* i [3] [172.16.254.5:102] [0] [32] [172.16.254.5]/17
      FC00:172:16:254::5          0    100    0 ?
* > i          FC00:172:16:254::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 ?
* > [5] [1:1] [0] [24] [10.1.102.0]/17
      0.0.0.0                      0    32768 ?
* > [5] [1:1] [0] [64] [FD00:10:1:101::]/29
      ::                          0    32768 ?
* > [5] [1:1] [0] [64] [FD00:10:1:102::]/29
      ::                          0    32768 ?

```

Leaf-01# show l2vpn evpn mac ip

IP Address	EVI	VLAN	MAC Address	Next Hop(s)
10.1.101.11	101	101	f4cf.e243.34c1	Gi1/0/4:101
FE80::F6CF:E2FF:FE43:34C1	101	101	f4cf.e243.34c1	Gi1/0/4:101
10.1.102.12	102	102	44d3.ca28.6cc2	FC00:172:16:254::4
10.1.102.13	102	102	44d3.ca28.6cc4	FC00:172:16:254::5
FE80::F6CF:E2FF:FE43:34C2	102	102	f4cf.e243.34c2	Gi1/0/5:102

Leaf-01# show l2route evpn mac ip

EVI	ETag	Prod	Mac Address	Next Hop(s)	Host IP

```

-----
101          0 L2VPN 10b3.d56a.8fc1                10.1.101.1
              V1101:0
101          0 L2VPN 10b3.d56a.8fc1                FD00:10:1:101::1
              V1101:0
101          0 BGP 7c21.0dbd.2741                10.1.101.1
V:10101 FC00:172:16:254::5
101          0 BGP 7c21.0dbd.2741                FD00:10:1:101::1
V:10101 FC00:172:16:254::5
101          0 BGP 7c21.0dbd.9541                10.1.101.1
V:10101 FC00:172:16:254::4
101          0 BGP 7c21.0dbd.9541                FD00:10:1:101::1
V:10101 FC00:172:16:254::4
101          0 L2VPN f4cf.e243.34c1                10.1.101.11
              Gil/0/4:101
101          0 L2VPN f4cf.e243.34c1                FE80::F6CF:E2FF:FE43:34C1
              Gil/0/4:101
102          0 L2VPN 10b3.d56a.8fcd                10.1.102.1
              V1102:0
102          0 L2VPN 10b3.d56a.8fcd                FD00:10:1:102::1
              V1102:0
102          0 BGP 44d3.ca28.6cc2                10.1.102.12
V:10102 FC00:172:16:254::4
102          0 BGP 44d3.ca28.6cc4                10.1.102.13
V:10102 FC00:172:16:254::5
102          0 BGP 7c21.0dbd.274d                10.1.102.1
V:10102 FC00:172:16:254::5
102          0 BGP 7c21.0dbd.274d                FD00:10:1:102::1
V:10102 FC00:172:16:254::5
102          0 BGP 7c21.0dbd.954d                10.1.102.1
V:10102 FC00:172:16:254::4
102          0 BGP 7c21.0dbd.954d                FD00:10:1:102::1
V:10102 FC00:172:16:254::4
102          0 L2VPN f4cf.e243.34c2                FE80::F6CF:E2FF:FE43:34C2
              Gil/0/5:102

```

To return to the configuration example, click [Example: Configuring BGP EVPN VXLAN with IPv6 Underlay, on page 740](#).

Outputs to Verify Configuration on VTEP 2

```

Leaf-02# show ipv6 route
IPv6 Routing Table - default - 20 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, ls - LISP destinations-summary
O FC00:172:16:13::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
C FC00:172:16:14::/64 [0/0]
  via GigabitEthernet1/0/1, directly connected
L FC00:172:16:14::4/128 [0/0]
  via GigabitEthernet1/0/1, receive
O FC00:172:16:15::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
O FC00:172:16:23::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
C FC00:172:16:24::/64 [0/0]
  via GigabitEthernet1/0/2, directly connected

```



```

L   FC00:172:16:24::4/128 [0/0]
    via GigabitEthernet1/0/2, receive
O   FC00:172:16:25::/64 [110/2]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
O   FC00:172:16:254::1/128 [110/1]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
O   FC00:172:16:254::2/128 [110/1]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
O   FC00:172:16:254::3/128 [110/2]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
LC  FC00:172:16:254::4/128 [0/0]
    via Loopback1, receive
O   FC00:172:16:254::5/128 [110/2]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
O   FC00:172:16:255::1/128 [110/1]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
O   FC00:172:16:255::2/128 [110/1]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
O   FC00:172:16:255::3/128 [110/2]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
LC  FC00:172:16:255::4/128 [0/0]
    via Loopback0, receive
O   FC00:172:16:255::5/128 [110/2]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
O   FC00:172:16:255::255/128 [110/1]
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FE92:B2D6, GigabitEthernet1/0/2
L   FF00::/8 [0/0]
    via Null0, receive

```

Leaf-02# **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

```

(*, FF0E:225::101), 00:05:12/never, RP FC00:172:16:255::255, flags: SCJ
  Incoming interface: GigabitEthernet1/0/2
  RPF nbr: FE80::7E21:DFE:FE92:B2D6
  Immediate Outgoing interface list:
    Tunnell1, Forward, 00:05:12/never

```

```

(FC00:172:16:254::3, FF0E:225::101), 00:03:23/00:01:40, flags: SJT
  Incoming interface: GigabitEthernet1/0/1
  RPF nbr: FE80::7E21:DFE:FE92:B2D6
  Inherited Outgoing interface list:
    Tunnell1, Forward, 00:05:12/never

```

```

(FC00:172:16:254::4, FF0E:225::101), 00:04:40/00:02:41, flags: SFJT
  Incoming interface: Loopback1
  RPF nbr: FE80::7E21:DFE:FE92:9500
  Immediate Outgoing interface list:

```

Verifying BGP EVPN VXLAN with IPv6 Underlay Configuration

```
GigabitEthernet1/0/1, Forward, 00:04:40/00:02:50
Inherited Outgoing interface list:
  Tunnell, Forward, 00:05:12/never
```

```
(FC00:172:16:254::5, FF0E:225::101), 00:04:33/00:00:01, flags: SJT
Incoming interface: GigabitEthernet1/0/1
RPF nbr: FE80::7E21:DFE:FE92:B2D6
Inherited Outgoing interface list:
  Tunnell, Forward, 00:05:12/never
```

Leaf-02# show nve peer

```
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag
```

Interface	VNI	Type	Peer-IP	RMAC/Num_RT	eVNI	state	flags	UP	time
nve1	50901	L3CP	FC00:172:16:254::3 \	10b3.d56a.8fc8	50901	UP	A/-/4	00:03:41	
nve1	50901	L3CP	FC00:172:16:254::5 \	7c21.0dbd.2748	50901	UP	A/-/4	00:03:41	
nve1	50901	L3CP	FC00:172:16:254::3 \	10b3.d56a.8fc8	50901	UP	A/M/6	00:03:41	
nve1	50901	L3CP	FC00:172:16:254::5 \	7c21.0dbd.2748	50901	UP	A/M/6	00:03:41	
nve1	10101	L2CP	FC00:172:16:254::3 \	6	10101	UP	N/A	00:03:41	
nve1	10101	L2CP	FC00:172:16:254::5 \	4	10101	UP	N/A	00:03:41	
nve1	10102	L2CP	FC00:172:16:254::3 \	6	10102	UP	N/A	00:03:41	
nve1	10102	L2CP	FC00:172:16:254::5 \	6	10102	UP	N/A	00:03:41	

Leaf-02# show bgp l2vpn evpn summary

```
BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 53, main routing table version 53
48 network entries using 18432 bytes of memory
74 path entries using 17168 bytes of memory
19/18 BGP path/bestpath attribute entries using 5624 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
15 BGP extended community entries using 784 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 42168 total bytes of memory
BGP activity 62/2 prefixes, 113/19 paths, scan interval 60 secs
48 networks peaked at 16:39:55 Oct 31 2022 UTC (00:01:46.413 ago)
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
FC00:172:16:255::1	4	65001	27	14	53	0	0	00:04:49	22
FC00:172:16:255::2	4	65001	27	14	53	0	0	00:04:39	22

Leaf-02# show bgp l2vpn evpn

```
BGP table version is 53, 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
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:101] [0] [48] [10B3D56A8FC1] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [0] [*]/20
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [32] [10.1.101.11]/24
    FC00:172:16:254::3
    0 100 0 ?
* i FC00:172:16:254::3
    0 100 0 ?
*>i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
    FC00:172:16:254::3
    0 100 0 ?
* i FC00:172:16:254::3
    0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::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
    FC00:172:16:254::3
    0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [10B3D56A8FC1] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::3
    0 100 0 ?
*> [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [0] [*]/20
    :: 32768 ?
*>i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC3] [0] [*]/20
    FC00:172:16:254::5
    0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [7C210DBD2741] [32] [10.1.101.1]/24
    FC00:172:16:254::5
    0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::5
    0 100 0 ?

```

```

*> [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
      :: 32768 ?
*> [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
      :: 32768 ?
*>i [2] [172.16.254.4:101] [0] [48] [F4CFE24334C1] [0] [*]/20
      FC00:172:16:254::3
      0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [F4CFE24334C1] [32] [10.1.101.11]/24
      FC00:172:16:254::3
      0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
      FC00:172:16:254::3
      0 100 0 ?
Route Distinguisher: 172.16.254.4:102
*>i [2] [172.16.254.4:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24
      FC00:172:16:254::3
      0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
      FC00:172:16:254::3
      0 100 0 ?
*> [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [0] [*]/20
      :: 32768 ?
*> [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
      :: 32768 ?
*>i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC4] [0] [*]/20
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
      FC00:172:16:254::5
      0 100 0 ?
*> [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] [FD00:10:1:102::1]/36
      :: 32768 ?
*>i [2] [172.16.254.4:102] [0] [48] [F4CFE24334C2] [0] [*]/20
      FC00:172:16:254::3
      0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36
      FC00:172:16:254::3
      0 100 0 ?
Route Distinguisher: 172.16.254.5:101
* i [2] [172.16.254.5:101] [0] [48] [44D3CA286CC3] [0] [*]/20
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [32] [10.1.101.1]/24
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
      FC00:172:16:254::5
      0 100 0 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
      FC00:172:16:254::5
      0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [0] [*]/20

```

```

FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
FC00:172:16:254::3
0 100 0 ?
*>i FC00:172:16:254::3
0 100 0 ?
Route Distinguisher: 172.16.254.4:102
*>i [3] [172.16.254.4:102] [0] [32] [172.16.254.3]/17
FC00:172:16:254::3
0 100 0 ?
*> [3] [172.16.254.4:102] [0] [32] [172.16.254.4]/17
:: 32768 ?
*>i [3] [172.16.254.4:102] [0] [32] [172.16.254.5]/17
FC00:172:16:254::5
0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [3] [172.16.254.5:102] [0] [32] [172.16.254.5]/17
FC00:172:16:254::5
0 100 0 ?
*>i FC00:172:16:254::5
0 100 0 ?
Route Distinguisher: 1:1 (default for vrf green)
* i [5] [1:1] [0] [24] [10.1.101.0]/17
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::3
0 100 0 ?
*> 0.0.0.0 32768 ?
* i [5] [1:1] [0] [24] [10.1.102.0]/17
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::3
0 100 0 ?
*> 0.0.0.0 32768 ?
* i [5] [1:1] [0] [64] [FD00:10:1:101::]/29
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::3
0 100 0 ?
*> :: 32768 ?
* i [5] [1:1] [0] [64] [FD00:10:1:102::]/29
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::3
0 100 0 ?

```

```
*>                               ::                               0                               32768 ?
```

```
Leaf-02# show l2vpn evpn mac ip
IP Address                               EVI    VLAN  MAC Address      Next Hop(s)
-----
10.1.101.11                             101    101   f4cf.e243.34c1  FC00:172:16:254::3
FE80::F6CF:E2FF:FE43:34C1               101    101   f4cf.e243.34c1  FC00:172:16:254::3
10.1.102.12                             102    102   44d3.ca28.6cc2  Gi1/0/12:102
10.1.102.13                             102    102   44d3.ca28.6cc4  FC00:172:16:254::5
FE80::F6CF:E2FF:FE43:34C2               102    102   f4cf.e243.34c2  FC00:172:16:254::3
```

```
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 FC00:172:16:254::3
101    0     BGP  10b3.d56a.8fc1   FD00:10:1:101::1
V:10101 FC00:172:16:254::3
101    0     BGP  7c21.0dbd.2741  10.1.101.1
V:10101 FC00:172:16:254::5
101    0     BGP  7c21.0dbd.2741  FD00:10:1:101::1
V:10101 FC00:172:16:254::5
101    0     L2VPN 7c21.0dbd.9541  10.1.101.1
V1101:0
101    0     L2VPN 7c21.0dbd.9541  FD00:10:1:101::1
V1101:0
101    0     BGP  f4cf.e243.34c1  10.1.101.11
V:10101 FC00:172:16:254::3
101    0     BGP  f4cf.e243.34c1  FE80::F6CF:E2FF:FE43:34C1
V:10101 FC00:172:16:254::3
102    0     BGP  10b3.d56a.8fcd  10.1.102.1
V:10102 FC00:172:16:254::3
102    0     BGP  10b3.d56a.8fcd  FD00:10:1:102::1
V:10102 FC00:172:16:254::3
102    0     L2VPN 44d3.ca28.6cc2  10.1.102.12
Gi1/0/12:102
102    0     BGP  44d3.ca28.6cc4  10.1.102.13
V:10102 FC00:172:16:254::5
102    0     BGP  7c21.0dbd.274d  10.1.102.1
V:10102 FC00:172:16:254::5
102    0     BGP  7c21.0dbd.274d  FD00:10:1:102::1
V:10102 FC00:172:16:254::5
102    0     L2VPN 7c21.0dbd.954d  10.1.102.1
V1102:0
102    0     L2VPN 7c21.0dbd.954d  FD00:10:1:102::1
V1102:0
102    0     BGP  f4cf.e243.34c2  FE80::F6CF:E2FF:FE43:34C2
V:10102 FC00:172:16:254::3
```

To return to the configuration example, click [Example: Configuring BGP EVPN VXLAN with IPv6 Underlay, on page 740](#).

Outputs to Verify Configuration on VTEP 3

```
Leaf-03# show ipv6 route
IPv6 Routing Table - default - 20 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
```

```

I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
Ndr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
ld - LISP dyn-eid, la - LISP away, le - LISP extranet-policy
lp - LISP publications, ls - LISP destinations-summary
O FC00:172:16:13::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
O FC00:172:16:14::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
C FC00:172:16:15::/64 [0/0]
  via GigabitEthernet1/0/1, directly connected
L FC00:172:16:15::5/128 [0/0]
  via GigabitEthernet1/0/1, receive
O FC00:172:16:23::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O FC00:172:16:24::/64 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
C FC00:172:16:25::/64 [0/0]
  via GigabitEthernet1/0/2, directly connected
L FC00:172:16:25::5/128 [0/0]
  via GigabitEthernet1/0/2, receive
O FC00:172:16:254::1/128 [110/1]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
O FC00:172:16:254::2/128 [110/1]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O FC00:172:16:254::3/128 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O FC00:172:16:254::4/128 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
LC FC00:172:16:254::5/128 [0/0]
  via Loopback1, receive
O FC00:172:16:255::1/128 [110/1]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
O FC00:172:16:255::2/128 [110/1]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O FC00:172:16:255::3/128 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O FC00:172:16:255::4/128 [110/2]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
LC FC00:172:16:255::5/128 [0/0]
  via Loopback0, receive
O FC00:172:16:255::255/128 [110/1]
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
  via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
L FF00::/8 [0/0]
  via Null0, receive

```

Leaf-03# **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

(*, FF0E:225::101), 00:05:13/never, RP FC00:172:16:255::255, flags: SCJ
  Incoming interface: GigabitEthernet1/0/1
  RPF nbr: FE80::7E21:DFE:FE92:B2D8
  Immediate Outgoing interface list:
    Tunnel1, Forward, 00:05:13/never

(FC00:172:16:254::3, FF0E:225::101), 00:04:55/00:02:48, flags: SJT
  Incoming interface: GigabitEthernet1/0/1
  RPF nbr: FE80::7E21:DFE:FE92:B2D8
  Inherited Outgoing interface list:
    Tunnel1, Forward, 00:05:13/never

(FC00:172:16:254::4, FF0E:225::101), 00:04:48/00:02:24, flags: SJT
  Incoming interface: GigabitEthernet1/0/1
  RPF nbr: FE80::7E21:DFE:FE92:B2D8
  Inherited Outgoing interface list:
    Tunnel1, Forward, 00:05:13/never

(FC00:172:16:254::5, FF0E:225::101), 00:04:41/00:02:50, flags: SFT
  Incoming interface: Loopback1
  RPF nbr: FE80::7E21:DFE:FE92:B2D8
  Immediate Outgoing interface list:
    GigabitEthernet1/0/2, Forward, 00:04:41/00:02:50
  Inherited Outgoing interface list:
    Tunnel1, Forward, 00:05:13/never

```

Leaf-03# **show nve peer**

'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag

Interface	VNI	Type	Peer-IP	RMAC/Num_RT	eVNI	state	flags	UP	time
nve1	50901	L3CP	FC00:172:16:254::3 \	10b3.d56a.8fc8	50901	UP	A/-/4	00:03:49	
nve1	50901	L3CP	FC00:172:16:254::4 \	7c21.0dbd.9548	50901	UP	A/-/4	00:03:49	
nve1	50901	L3CP	FC00:172:16:254::3 \	10b3.d56a.8fc8	50901	UP	A/M/6	00:03:49	
nve1	50901	L3CP	FC00:172:16:254::4 \	7c21.0dbd.9548	50901	UP	A/M/6	00:03:49	
nve1	10101	L2CP	FC00:172:16:254::3 \	6	10101	UP	N/A	00:03:49	
nve1	10101	L2CP	FC00:172:16:254::4 \	4	10101	UP	N/A	00:03:49	
nve1	10102	L2CP	FC00:172:16:254::3 \	6	10102	UP	N/A	00:03:49	
nve1	10102	L2CP	FC00:172:16:254::4 \	6	10102	UP	N/A	00:03:49	

Leaf-03# **show bgp l2vpn evpn summary**

```

BGP router identifier 172.16.255.5, local AS number 65001
BGP table version is 53, main routing table version 53
48 network entries using 18432 bytes of memory
74 path entries using 17168 bytes of memory
19/18 BGP path/bestpath attribute entries using 5624 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
15 BGP extended community entries using 784 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 42168 total bytes of memory
BGP activity 60/0 prefixes, 100/6 paths, scan interval 60 secs

```


48 networks peaked at 16:39:55 Oct 31 2022 UTC (00:01:54.313 ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
FC00:172:16:255::1	4	65001	28	13	53	0	0	00:04:56	22
FC00:172:16:255::2	4	65001	28	14	53	0	0	00:04:49	22

```
Leaf-03#show bgp l2vpn evpn
BGP table version is 53, 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: 172.16.254.3:101					
* i [2] [172.16.254.3:101] [0] [48] [10B3D56A8FC1] [32] [10.1.101.1]/24	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:101] [0] [48] [10B3D56A8FC1] [128] [FD00:10:1:101::1]/36	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [0] [*]/20	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
*>i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [32] [10.1.101.11]/24	FC00:172:16:254::3				
		0	100	0	?
* i	FC00:172:16:254::3				
		0	100	0	?
*>i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36	FC00:172:16:254::3				
		0	100	0	?
* i	FC00:172:16:254::3				
		0	100	0	?
Route Distinguisher: 172.16.254.3:102					
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				

```

                                0    100    0 ?
Route Distinguisher: 172.16.254.4:101
* i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [0] [*]/20
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
Route Distinguisher: 172.16.254.4:102
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [0] [*]/20
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [32] [10.1.102.1]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [128] [FD00:10:1:102::1]/36
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
Route Distinguisher: 172.16.254.5:101
*>i [2] [172.16.254.5:101] [0] [48] [10B3D56A8FC1] [32] [10.1.101.1]/24
    FC00:172:16:254::3
                                0    100    0 ?
*>i [2] [172.16.254.5:101] [0] [48] [10B3D56A8FC1] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::3
                                0    100    0 ?
*>i [2] [172.16.254.5:101] [0] [48] [44D3CA286CC1] [0] [*]/20
    FC00:172:16:254::4
                                0    100    0 ?
*> [2] [172.16.254.5:101] [0] [48] [44D3CA286CC3] [0] [*]/20
    ::                               32768 ?
*> [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [32] [10.1.101.1]/24
    ::                               32768 ?
*> [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
    ::                               32768 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::4
                                0    100    0 ?
*>i [2] [172.16.254.5:101] [0] [48] [F4CFE24334C1] [0] [*]/20
    FC00:172:16:254::3
                                0    100    0 ?
*>i [2] [172.16.254.5:101] [0] [48] [F4CFE24334C1] [32] [10.1.101.11]/24

```

```

FC00:172:16:254::3
0 100 0 ?
*>i [2] [172.16.254.5:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
FC00:172:16:254::3
0 100 0 ?
Route Distinguisher: 172.16.254.5:102
*>i [2] [172.16.254.5:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24
FC00:172:16:254::3
0 100 0 ?
*>i [2] [172.16.254.5:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
FC00:172:16:254::3
0 100 0 ?
*>i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC2] [0] [*]/20
FC00:172:16:254::4
0 100 0 ?
*>i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
FC00:172:16:254::4
0 100 0 ?
*> [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [0] [*]/20
:: 32768 ?
*> [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
:: 32768 ?
*> [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
:: 32768 ?
*> [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
:: 32768 ?
*>i [2] [172.16.254.5:102] [0] [48] [7C210DBD954D] [32] [10.1.102.1]/24
FC00:172:16:254::4
0 100 0 ?
*>i [2] [172.16.254.5:102] [0] [48] [7C210DBD954D] [128] [FD00:10:1:102::1]/36
FC00:172:16:254::4
0 100 0 ?
*>i [2] [172.16.254.5:102] [0] [48] [F4CFE24334C2] [0] [*]/20
FC00:172:16:254::3
0 100 0 ?
*>i [2] [172.16.254.5:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36
FC00:172:16:254::3
0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
FC00:172:16:254::3
0 100 0 ?
*>i FC00:172:16:254::3
0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [3] [172.16.254.4:102] [0] [32] [172.16.254.4]/17
FC00:172:16:254::4
0 100 0 ?
*>i FC00:172:16:254::4
0 100 0 ?
Route Distinguisher: 172.16.254.5:102
*>i [3] [172.16.254.5:102] [0] [32] [172.16.254.3]/17
FC00:172:16:254::3
0 100 0 ?
*>i [3] [172.16.254.5:102] [0] [32] [172.16.254.4]/17
FC00:172:16:254::4
0 100 0 ?
*> [3] [172.16.254.5:102] [0] [32] [172.16.254.5]/17
:: 32768 ?
Route Distinguisher: 1:1 (default for vrf green)
* i [5] [1:1] [0] [24] [10.1.101.0]/17
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::3

```

Verifying BGP EVPN VXLAN with IPv6 Underlay Configuration

```

* >
* i [5][1:1][0][24][10.1.102.0]/17
      FC00:172:16:254::3
      0 100 0 ?
* i
      FC00:172:16:254::3
      0 100 0 ?
* >
* i [5][1:1][0][64][FD00:10:1:101:]/29
      FC00:172:16:254::3
      0 100 0 ?
* i
      FC00:172:16:254::3
      0 100 0 ?
* >
* i [5][1:1][0][64][FD00:10:1:102:]/29
      FC00:172:16:254::3
      0 100 0 ?
* i
      FC00:172:16:254::3
      0 100 0 ?
* >
      ::
      0 100 0 ?
* i [5][1:1][0][64][FD00:10:1:102:]/29
      FC00:172:16:254::3
      0 100 0 ?
* i
      FC00:172:16:254::3
      0 100 0 ?
* >
      ::
      0 100 0 ?

```

Leaf-03# show l2vpn evpn mac ip

IP Address	EVI	VLAN	MAC Address	Next Hop(s)
10.1.101.11	101	101	f4cf.e243.34c1	FC00:172:16:254::3
FE80::F6CF:E2FF:FE43:34C1	101	101	f4cf.e243.34c1	FC00:172:16:254::3
10.1.102.12	102	102	44d3.ca28.6cc2	FC00:172:16:254::4
10.1.102.13	102	102	44d3.ca28.6cc4	Gi1/0/17:102
FE80::F6CF:E2FF:FE43:34C2	102	102	f4cf.e243.34c2	FC00:172:16:254::3

Leaf-03# show l2route evpn mac ip

EVI	Etag	Prod	Mac Address	Next Hop(s)	Host IP
101	0	BGP	10b3.d56a.8fc1		10.1.101.1
		V:10101	FC00:172:16:254::3		
101	0	BGP	10b3.d56a.8fc1		FD00:10:1:101::1
		V:10101	FC00:172:16:254::3		
101	0	L2VPN	7c21.0dbd.2741		10.1.101.1
			Vl1101:0		
101	0	L2VPN	7c21.0dbd.2741		FD00:10:1:101::1
			Vl1101:0		
101	0	BGP	7c21.0dbd.9541		10.1.101.1
		V:10101	FC00:172:16:254::4		
101	0	BGP	7c21.0dbd.9541		FD00:10:1:101::1
		V:10101	FC00:172:16:254::4		
101	0	BGP	f4cf.e243.34c1		10.1.101.11
		V:10101	FC00:172:16:254::3		
101	0	BGP	f4cf.e243.34c1		FE80::F6CF:E2FF:FE43:34C1
		V:10101	FC00:172:16:254::3		
102	0	BGP	10b3.d56a.8fcd		10.1.102.1
		V:10102	FC00:172:16:254::3		
102	0	BGP	10b3.d56a.8fcd		FD00:10:1:102::1
		V:10102	FC00:172:16:254::3		
102	0	BGP	44d3.ca28.6cc2		10.1.102.12
		V:10102	FC00:172:16:254::4		
102	0	L2VPN	44d3.ca28.6cc4		10.1.102.13
			Gi1/0/17:102		
102	0	L2VPN	7c21.0dbd.274d		10.1.102.1
			Vl1102:0		

```

102          0 L2VPN 7c21.0dbd.274d          FD00:10:1:102::1
              V1102:0
102          0 BGP 7c21.0dbd.954d          10.1.102.1
              V:10102 FC00:172:16:254::4
102          0 BGP 7c21.0dbd.954d          FD00:10:1:102::1
              V:10102 FC00:172:16:254::4
102          0 BGP f4cf.e243.34c2          FE80::F6CF:E2FF:FE43:34C2
              V:10102 FC00:172:16:254::3

```

To return to the configuration example, click [Example: Configuring BGP EVPN VXLAN with IPv6 Underlay](#), on page 740.

Outputs to Verify Configuration on Spine 1

```

Spine-01# show ipv6 route
IPv6 Routing Table - default - 21 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
        B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
        I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
        EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
        NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
        OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
        ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
        ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
        lp - LISP publications, ls - LISP destinations-summary
C   FC00:172:16:13::/64 [0/0]
    via GigabitEthernet1/0/1, directly connected
L   FC00:172:16:13::1/128 [0/0]
    via GigabitEthernet1/0/1, receive
C   FC00:172:16:14::/64 [0/0]
    via GigabitEthernet1/0/2, directly connected
L   FC00:172:16:14::1/128 [0/0]
    via GigabitEthernet1/0/2, receive
C   FC00:172:16:15::/64 [0/0]
    via GigabitEthernet1/0/3, directly connected
L   FC00:172:16:15::1/128 [0/0]
    via GigabitEthernet1/0/3, receive
O   FC00:172:16:23::/64 [110/2]
    via FE80::12B3:D5FF:FE6A:8FE4, GigabitEthernet1/0/1
O   FC00:172:16:24::/64 [110/2]
    via FE80::7E21:DFE:FEBD:9564, GigabitEthernet1/0/2
O   FC00:172:16:25::/64 [110/2]
    via FE80::7E21:DFE:FEBD:2764, GigabitEthernet1/0/3
LC  FC00:172:16:254::1/128 [0/0]
    via Loopback1, receive
O   FC00:172:16:254::2/128 [110/2]
    via FE80::7E21:DFE:FEBD:9564, GigabitEthernet1/0/2
    via FE80::12B3:D5FF:FE6A:8FE4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FEBD:2764, GigabitEthernet1/0/3
O   FC00:172:16:254::3/128 [110/1]
    via FE80::12B3:D5FF:FE6A:8FE4, GigabitEthernet1/0/1
O   FC00:172:16:254::4/128 [110/1]
    via FE80::7E21:DFE:FEBD:9564, GigabitEthernet1/0/2
O   FC00:172:16:254::5/128 [110/1]
    via FE80::7E21:DFE:FEBD:2764, GigabitEthernet1/0/3
LC  FC00:172:16:255::1/128 [0/0]
    via Loopback0, receive
O   FC00:172:16:255::2/128 [110/2]
    via FE80::7E21:DFE:FEBD:9564, GigabitEthernet1/0/2
    via FE80::12B3:D5FF:FE6A:8FE4, GigabitEthernet1/0/1
    via FE80::7E21:DFE:FEBD:2764, GigabitEthernet1/0/3
O   FC00:172:16:255::3/128 [110/1]

```

```

    via FE80::12B3:D5FF:FE6A:8FE4, GigabitEthernet1/0/1
O  FC00:172:16:255::4/128 [110/1]
    via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
O  FC00:172:16:255::5/128 [110/1]
    via FE80::7E21:DFF:FEBD:2764, GigabitEthernet1/0/3
LC FC00:172:16:255::255/128 [0/0]
    via Loopback2, receive
L  FF00::/8 [0/0]
    via Null0, receive

```

Spine-01# **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

```

(*, FF0E:225::101), 00:05:14/00:03:17, RP FC00:172:16:255::255, flags: S
  Incoming interface: Tunnel2
  RPF nbr: FC00:172:16:255::255
  Immediate Outgoing interface list:
    GigabitEthernet1/0/3, Forward, 00:05:14/00:03:17

(FC00:172:16:254::3, FF0E:225::101), 00:05:14/00:01:13, flags: ST
  Incoming interface: GigabitEthernet1/0/1
  RPF nbr: FE80::12B3:D5FF:FE6A:8FE4
  Immediate Outgoing interface list:
    GigabitEthernet1/0/3, Forward, 00:05:14/00:03:17
    GigabitEthernet1/0/2, Forward, 00:03:49/00:02:40

(FC00:172:16:254::4, FF0E:225::101), 00:05:06/00:02:04, flags: ST
  Incoming interface: GigabitEthernet1/0/2
  RPF nbr: FE80::7E21:DFF:FEBD:9564
  Immediate Outgoing interface list:
    GigabitEthernet1/0/3, Forward, 00:05:06/00:03:27
    GigabitEthernet1/0/1, Forward, 00:05:06/00:03:23

(FC00:172:16:254::5, FF0E:225::101), 00:04:59/00:03:17, RP FC00:172:16:255::255, flags: SPR
  Incoming interface: Tunnel2
  RPF nbr: FC00:172:16:255::255
  Immediate Outgoing interface list:
    GigabitEthernet1/0/3, Null, 00:04:59/00:03:17

```

Spine-01# **show bgp l2vpn evpn summary**

```

BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 31, main routing table version 31
30 network entries using 11520 bytes of memory
68 path entries using 15776 bytes of memory
18/16 BGP path/bestpath attribute entries using 5328 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
15 BGP extended community entries using 800 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 33544 total bytes of memory
BGP activity 30/0 prefixes, 68/0 paths, scan interval 60 secs

```

30 networks peaked at 16:39:55 Oct 31 2022 UTC (00:02:12.681 ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
FC00:172:16:255::2	4	65001	28	28	31	0	0	00:05:03	30
FC00:172:16:255::3	4	65001	16	28	31	0	0	00:05:18	14
FC00:172:16:255::4	4	65001	14	28	31	0	0	00:05:15	12
FC00:172:16:255::5	4	65001	14	28	31	0	0	00:05:14	12

Spine-01# show bgp l2vpn evpn

BGP table version is 31, 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	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:101] [0] [48] [10B3D56A8FC1] [128] [FD00:10:1:101::1]/36	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [0] [*]/20	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [32] [10.1.101.11]/24	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
Route Distinguisher: 172.16.254.3:102					
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				
		0	100	0	?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20	FC00:172:16:254::3				
		0	100	0	?
*>i	FC00:172:16:254::3				

```

                                0    100    0 ?
* i [2][172.16.254.3:102][0][48][F4CFE24334C2][128][FE80::F6CF:E2FF:FE43:34C2]/36
    FC00:172:16:254::3
                                0    100    0 ?
*>i          FC00:172:16:254::3
                                0    100    0 ?
Route Distinguisher: 172.16.254.4:101
* i [2][172.16.254.4:101][0][48][44D3CA286CC1][0][*]/20
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2][172.16.254.4:101][0][48][7C210DBD9541][128][FD00:10:1:101::1]/36
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
Route Distinguisher: 172.16.254.4:102
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][0][*]/20
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][32][10.1.102.12]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
* i [2][172.16.254.4:102][0][48][7C210DBD954D][128][FD00:10:1:102::1]/36
    FC00:172:16:254::4
                                0    100    0 ?
*>i          FC00:172:16:254::4
                                0    100    0 ?
Route Distinguisher: 172.16.254.5:101
* i [2][172.16.254.5:101][0][48][44D3CA286CC3][0][*]/20
    FC00:172:16:254::5
                                0    100    0 ?
*>i          FC00:172:16:254::5
                                0    100    0 ?
* i [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
    FC00:172:16:254::5
                                0    100    0 ?
*>i          FC00:172:16:254::5
                                0    100    0 ?
* i [2][172.16.254.5:101][0][48][7C210DBD2741][128][FD00:10:1:101::1]/36
    FC00:172:16:254::5
                                0    100    0 ?
*>i          FC00:172:16:254::5
                                0    100    0 ?
Route Distinguisher: 172.16.254.5:102
* i [2][172.16.254.5:102][0][48][44D3CA286CC4][0][*]/20
    FC00:172:16:254::5
                                0    100    0 ?
*>i          FC00:172:16:254::5

```



```

0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
FC00:172:16:254::5
*>i FC00:172:16:254::5 0 100 0 ?
0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
FC00:172:16:254::5
*>i FC00:172:16:254::5 0 100 0 ?
0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
FC00:172:16:254::5
*>i FC00:172:16:254::5 0 100 0 ?
0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
FC00:172:16:254::3
*>i FC00:172:16:254::3 0 100 0 ?
0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [3] [172.16.254.4:102] [0] [32] [172.16.254.4]/17
FC00:172:16:254::4
*>i FC00:172:16:254::4 0 100 0 ?
0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [3] [172.16.254.5:102] [0] [32] [172.16.254.5]/17
FC00:172:16:254::5
*>i FC00:172:16:254::5 0 100 0 ?
0 100 0 ?
Route Distinguisher: 1:1
* i [5] [1:1] [0] [24] [10.1.101.0]/17
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::4 0 100 0 ?
*>i FC00:172:16:254::3 0 100 0 ?
* i FC00:172:16:254::5 0 100 0 ?
* i [5] [1:1] [0] [24] [10.1.102.0]/17
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::4 0 100 0 ?
*>i FC00:172:16:254::3 0 100 0 ?
* i FC00:172:16:254::5 0 100 0 ?
* i [5] [1:1] [0] [64] [FD00:10:1:101::]/29
FC00:172:16:254::3
0 100 0 ?
* i FC00:172:16:254::4 0 100 0 ?
*>i FC00:172:16:254::3 0 100 0 ?
* i FC00:172:16:254::5 0 100 0 ?
* i [5] [1:1] [0] [64] [FD00:10:1:102::]/29
FC00:172:16:254::3

```

```

* i          FC00:172:16:254::4          0    100    0 ?
*>i         FC00:172:16:254::3          0    100    0 ?
* i          FC00:172:16:254::5          0    100    0 ?

```

```

Spine-01# show ipv6 pim anycast-rp
Anycast RP Peers For FC00:172:16:255::255    Last Register/Register-Stop received
FC00:172:16:254::2 00:05:25/00:05:25

```

To return to the configuration example, click [Example: Configuring BGP EVPN VXLAN with IPv6 Underlay, on page 740](#).

Outputs to Verify Configuration on Spine 2

```

Spine-02# show ipv6 route
IPv6 Routing Table - default - 21 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy
       lp - LISP publications, ls - LISP destinations-summary
O FC00:172:16:13::/64 [110/2]
  via FE80::12B3:D5FF:FE6A:8FD6, GigabitEthernet1/0/1
O FC00:172:16:14::/64 [110/2]
  via FE80::7E21:DFF:FEBD:9556, GigabitEthernet1/0/2
O FC00:172:16:15::/64 [110/2]
  via FE80::7E21:DFF:FEBD:2756, GigabitEthernet1/0/3
C FC00:172:16:23::/64 [0/0]
  via GigabitEthernet1/0/1, directly connected
L FC00:172:16:23::2/128 [0/0]
  via GigabitEthernet1/0/1, receive
C FC00:172:16:24::/64 [0/0]
  via GigabitEthernet1/0/2, directly connected
L FC00:172:16:24::2/128 [0/0]
  via GigabitEthernet1/0/2, receive
C FC00:172:16:25::/64 [0/0]
  via GigabitEthernet1/0/3, directly connected
L FC00:172:16:25::2/128 [0/0]
  via GigabitEthernet1/0/3, receive
O FC00:172:16:254::1/128 [110/2]
  via FE80::7E21:DFF:FEBD:9556, GigabitEthernet1/0/2
  via FE80::7E21:DFF:FEBD:2756, GigabitEthernet1/0/3
  via FE80::12B3:D5FF:FE6A:8FD6, GigabitEthernet1/0/1
LC FC00:172:16:254::2/128 [0/0]
  via Loopback1, receive
O FC00:172:16:254::3/128 [110/1]
  via FE80::12B3:D5FF:FE6A:8FD6, GigabitEthernet1/0/1
O FC00:172:16:254::4/128 [110/1]
  via FE80::7E21:DFF:FEBD:9556, GigabitEthernet1/0/2
O FC00:172:16:254::5/128 [110/1]
  via FE80::7E21:DFF:FEBD:2756, GigabitEthernet1/0/3
O FC00:172:16:255::1/128 [110/2]
  via FE80::7E21:DFF:FEBD:9556, GigabitEthernet1/0/2
  via FE80::7E21:DFF:FEBD:2756, GigabitEthernet1/0/3

```

```

    via FE80::12B3:D5FF:FE6A:8FD6, GigabitEthernet1/0/1
LC FC00:172:16:255::2/128 [0/0]
    via Loopback0, receive
O FC00:172:16:255::3/128 [110/1]
    via FE80::12B3:D5FF:FE6A:8FD6, GigabitEthernet1/0/1
O FC00:172:16:255::4/128 [110/1]
    via FE80::7E21:DFE:FEBD:9556, GigabitEthernet1/0/2
O FC00:172:16:255::5/128 [110/1]
    via FE80::7E21:DFE:FEBD:2756, GigabitEthernet1/0/3
LC FC00:172:16:255::255/128 [0/0]
    via Loopback2, receive
L FF00::/8 [0/0]
    via Null0, receive

```

Spine-02# **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

```

(*, FF0E:225::101), 00:05:17/00:03:21, RP FC00:172:16:255::255, flags: S
Incoming interface: Tunnel2
RPF nbr: FC00:172:16:255::255
Immediate outgoing interface list:
  GigabitEthernet1/0/1, Forward, 00:05:17/00:03:14
  GigabitEthernet1/0/2, Forward, 00:05:15/00:03:21

(FC00:172:16:254::3, FF0E:225::101), 00:05:17/00:02:55, RP FC00:172:16:255::255, flags: SPR

Incoming interface: Tunnel2
RPF nbr: FC00:172:16:255::255
Immediate outgoing interface list:
  GigabitEthernet1/0/1, Null, 00:05:17/00:03:14
  GigabitEthernet1/0/2, Null, 00:02:24/00:03:21

(FC00:172:16:254::3, FF0E:225::101), 00:04:27/00:02:55, flags: S
Incoming interface: GigabitEthernet1/0/1
RPF nbr: FE80::12B3:D5FF:FE6A:8FD6
Inherited outgoing interface list:
  GigabitEthernet1/0/2, Forward, 00:05:15/00:03:21

(FC00:172:16:254::4, FF0E:225::101), 00:05:15/00:02:50, RP FC00:172:16:255::255, flags: SPR

Incoming interface: Tunnel2
RPF nbr: FC00:172:16:255::255
Immediate outgoing interface list:
  GigabitEthernet1/0/2, Null, 00:05:15/00:03:21
  GigabitEthernet1/0/1, Null, 00:05:15/00:03:14

(FC00:172:16:254::4, FF0E:225::101), 00:05:15/00:02:50, flags: S
Incoming interface: GigabitEthernet1/0/2
RPF nbr: FE80::7E21:DFE:FEBD:9556
Inherited outgoing interface list:
  GigabitEthernet1/0/1, Forward, 00:05:17/00:03:14

```

```
(FC00:172:16:254::5, FF0E:225::101), 00:05:08/00:02:01, flags: ST
Incoming interface: GigabitEthernet1/0/3
RPF nbr: FE80::7E21:DFE:FEBD:2756
Inherited Outgoing interface list:
  GigabitEthernet1/0/1, Forward, 00:05:17/00:03:14
  GigabitEthernet1/0/2, Forward, 00:05:15/00:03:21
```

```
Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 31, main routing table version 31
30 network entries using 11520 bytes of memory
68 path entries using 15776 bytes of memory
18/16 BGP path/bestpath attribute entries using 5328 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
15 BGP extended community entries using 800 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 33544 total bytes of memory
BGP activity 30/0 prefixes, 68/0 paths, scan interval 60 secs
30 networks peaked at 16:39:55 Oct 31 2022 UTC (00:02:21.575 ago)
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
FC00:172:16:255::1	4	65001	28	28	31	0	0	00:05:12	30
FC00:172:16:255::3	4	65001	16	28	31	0	0	00:05:16	14
FC00:172:16:255::4	4	65001	14	28	31	0	0	00:05:15	12
FC00:172:16:255::5	4	65001	14	28	31	0	0	00:05:16	12

```
Spine-02# show bgp l2vpn evpn
BGP table version is 31, 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	FC00:172:16:254::3		0	100	0 ?
**>i	FC00:172:16:254::3		0	100	0 ?
* i [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FD00:10:1:101::1]/36	FC00:172:16:254::3		0	100	0 ?
**>i	FC00:172:16:254::3		0	100	0 ?
* i [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20	FC00:172:16:254::3		0	100	0 ?
**>i	FC00:172:16:254::3		0	100	0 ?
* i [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24	FC00:172:16:254::3		0	100	0 ?
**>i	FC00:172:16:254::3		0	100	0 ?

```

* i [2] [172.16.254.3:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
Route Distinguisher: 172.16.254.4:101
* i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [0] [*]/20
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
* i [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
* i [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [0] [*]/20
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [32] [10.1.102.1]/24
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [128] [FD00:10:1:102::1]/36
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
Route Distinguisher: 172.16.254.5:101

```

```

* i [2][172.16.254.5:101][0][48][44D3CA286CC3][0][*]/20
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
* i [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
* i [2][172.16.254.5:101][0][48][7C210DBD2741][128][FD00:10:1:101::1]/36
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [2][172.16.254.5:102][0][48][44D3CA286CC4][0][*]/20
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
* i [2][172.16.254.5:102][0][48][44D3CA286CC4][32][10.1.102.13]/24
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
* i [2][172.16.254.5:102][0][48][7C210DBD274D][32][10.1.102.1]/24
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
* i [2][172.16.254.5:102][0][48][7C210DBD274D][128][FD00:10:1:102::1]/36
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3][172.16.254.3:102][0][32][172.16.254.3]/17
    FC00:172:16:254::3
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [3][172.16.254.4:102][0][32][172.16.254.4]/17
    FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::4
    0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [3][172.16.254.5:102][0][32][172.16.254.5]/17
    FC00:172:16:254::5
    0 100 0 ?
*>i FC00:172:16:254::5
    0 100 0 ?
Route Distinguisher: 1:1
* i [5][1:1][0][24][10.1.101.0]/17
    FC00:172:16:254::3
    0 100 0 ?
* i FC00:172:16:254::4
    0 100 0 ?
*>i FC00:172:16:254::3
    0 100 0 ?
* i FC00:172:16:254::5
    0 100 0 ?

```

```

* i [5] [1:1] [0] [24] [10.1.102.0]/17
    FC00:172:16:254::3          0 100 0 ?
* i          FC00:172:16:254::4          0 100 0 ?
*>i          FC00:172:16:254::3          0 100 0 ?
* i          FC00:172:16:254::5          0 100 0 ?
* i [5] [1:1] [0] [64] [FD00:10:1:101::]/29
    FC00:172:16:254::3          0 100 0 ?
* i          FC00:172:16:254::4          0 100 0 ?
*>i          FC00:172:16:254::3          0 100 0 ?
* i          FC00:172:16:254::5          0 100 0 ?
* i [5] [1:1] [0] [64] [FD00:10:1:102::]/29
    FC00:172:16:254::3          0 100 0 ?
* i          FC00:172:16:254::4          0 100 0 ?
*>i          FC00:172:16:254::3          0 100 0 ?
* i          FC00:172:16:254::5          0 100 0 ?

```

```

Spine-02# show ipv6 pim anycast-rp
Anycast RP Peers For FC00:172:16:255::255 Last Register/Register-Stop received
FC00:172:16:254::1 00:05:27/00:05:27


```

To return to the configuration example, click [Example: Configuring BGP EVPN VXLAN with IPv6 Underlay, on page 740](#).

Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay

This example shows how to configure a BGP EVPN VXLAN fabric with a dual stack (both IPv4 and IPv6) underlay. The following tables provide sample configuration of the VTEPs (Leaf-01, Leaf-02, Leaf-03) and the Spine nodes (Spine-01, Spine-02) in the topology [Figure 72: BGP EVPN VXLAN with a Dual Stack Underlay](#).

Figure 72: BGP EVPN VXLAN with a Dual Stack Underlay



470878

Configure the VTEPs

VTEP 1	VTEP 2	VTEP 3
--------	--------	--------

VTEP 1	VTEP 2	VTEP 3
<pre>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 ! ipv6 unicast-routing ! 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 replication-type ingress ! 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 ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3</pre>	<pre>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 multicast-routing ! 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 replication-type ingress ! 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 ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 ipv6 address</pre>	<pre>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 ! ipv6 unicast-routing ipv6 multicast-routing ! 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 replication-type ingress ! 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 172.16.255.5 255.255.255.255 ip ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.3 ipv6 address</pre>

VTEP 1	VTEP 2	VTEP 3
<pre> 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 GigabitEthernet1/0/4 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/5 switchport access vlan 102 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address FD00:10:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 ipv6 address FD00:10:1:102::1/64 ipv6 enable ! 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 225.0.0.101 member vni 10102 ingress-replication member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.3 </pre>	<pre> FC00:172:16:255::4/128 ipv6 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 ipv6 address FC00:172:16:254::4/128 ipv6 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 ipv6 address FC00:172:16:14::4/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! 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 ipv6 address FC00:172:16:24::4/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/11 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/12 switchport access vlan 102 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address FD00:10:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 </pre>	<pre> FC00:172:16:255::5/128 ipv6 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 ipv6 address FC00:172:16:254::5/128 ipv6 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 ipv6 address FC00:172:16:15::5/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! 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 ipv6 address FC00:172:16:25::5/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! interface GigabitEthernet1/0/16 switchport access vlan 101 switchport mode access ! interface GigabitEthernet1/0/17 switchport access vlan 102 switchport mode access ! interface Vlan101 vrf forwarding green ip address 10.1.101.1 255.255.255.0 ipv6 address FD00:10:1:101::1/64 ipv6 enable ! interface Vlan102 vrf forwarding green ip address 10.1.102.1 255.255.255.0 </pre>

Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay

VTEP 1	VTEP 2	VTEP 3
<pre> ! router bgp 65001 bgp router-id 172.16.255.3 bgp log-neighbor-changes 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 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 exit-address-family ! address-family ipv6 vrf green redistribute connected advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! end </pre>	<pre> ipv6 address FD00:10:1:102::1/64 ipv6 enable ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4 member vni 10101 mcast-group 225.0.0.101 FF0E:225::101 member vni 10102 ingress-replication member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.4 ! router bgp 65001 bgp router-id 172.16.255.4 bgp log-neighbor-changes 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 neighbor FC00:172:16:255::1 remote-as 65001 neighbor FC00:172:16:255::1 update-source Loopback0 neighbor FC00:172:16:255::2 remote-as 65001 neighbor FC00:172:16:255::2 update-source Loopback0 ! 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 neighbor FC00:172:16:255::1 activate neighbor FC00:172:16:255::1 send-community both neighbor FC00:172:16:255::2 activate neighbor FC00:172:16:255::2 send-community both exit-address-family </pre>	<pre> ipv6 address FD00:10:1:102::1/64 ipv6 enable ! interface Vlan901 vrf forwarding green ip unnumbered Loopback1 ipv6 enable no autostate ! interface nve1 no ip address source-interface Loopback1 host-reachability protocol bgp vxlan encapsulation dual-stack prefer-ipv6 underlay-mcast ipv4 member vni 10101 mcast-group 225.0.0.101 FF0E:225::101 member vni 10102 ingress-replication member vni 50901 vrf green ! router ospf 1 router-id 172.16.255.5 ! router bgp 65001 bgp router-id 172.16.255.5 bgp log-neighbor-changes neighbor FC00:172:16:255::1 remote-as 65001 neighbor FC00:172:16:255::1 update-source Loopback0 neighbor FC00:172:16:255::2 remote-as 65001 neighbor FC00:172:16:255::2 update-source Loopback0 ! address-family l2vpn evpn neighbor FC00:172:16:255::1 activate neighbor FC00:172:16:255::1 send-community both neighbor FC00:172:16:255::2 activate neighbor FC00:172:16:255::2 send-community both exit-address-family ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected exit-address-family ! address-family ipv6 vrf green redistribute connected advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! </pre>

VTEP 1	VTEP 2	VTEP 3
	<pre> ! address-family ipv4 vrf green advertise l2vpn evpn redistribute connected exit-address-family ! address-family ipv6 vrf green redistribute connected advertise l2vpn evpn exit-address-family ! ip pim rp-address 172.16.255.255 ! ipv6 router ospf 1 router-id 172.16.255.4 ! ipv6 pim rp-address FC00:172:16:255::255 ! end </pre>	<pre> ipv6 router ospf 1 router-id 172.16.255.5 ! ipv6 pim rp-address FC00:172:16:255::255 ! ! end </pre>

Table 65: Configure the Spine Nodes

Spine 1	Spine 2
---------	---------

Spine 1	Spine 2
<pre> Spine-01# show running-config hostname Spine-01 ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ipv6 multicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ipv6 address FC00:172:16:255::1/128 ipv6 ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.1 255.255.255.255 ip ospf 1 area 0 ipv6 address FC00:172:16:254::1/128 ipv6 ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip ospf 1 area 0 ipv6 address FC00:172:16:255::255/128 ipv6 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 ipv6 address FC00:172:16:13::1/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! 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 ipv6 address FC00:172:16:14::1/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! 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 ipv6 address FC00:172:16:15::1/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point </pre>	<pre> Spine-02# show running-config hostname Spine-02 ! ip routing ! ip multicast-routing ! ipv6 unicast-routing ipv6 multicast-routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.2 255.255.255.255 ip ospf 1 area 0 ipv6 address FC00:172:16:255::2/128 ipv6 ospf 1 area 0 ! interface Loopback1 ip address 172.16.254.2 255.255.255.255 ip ospf 1 area 0 ipv6 address FC00:172:16:254::2/128 ipv6 ospf 1 area 0 ! interface Loopback2 ip address 172.16.255.255 255.255.255.255 ip ospf 1 area 0 ipv6 address FC00:172:16:255::255/128 ipv6 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 ipv6 address FC00:172:16:23::2/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! 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 ipv6 address FC00:172:16:24::2/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point ! 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 ipv6 address FC00:172:16:25::2/64 ipv6 enable ipv6 ospf 1 area 0 ipv6 ospf network point-to-point </pre>

Spine 1	Spine 2
<pre> ! router ospf 1 router-id 172.16.255.1 ! router bgp 65001 bgp router-id 172.16.255.1 bgp log-neighbor-changes 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 FC00:172:16:255::2 remote-as 65001 neighbor FC00:172:16:255::2 update-source Loopback0 neighbor FC00:172:16:255::4 remote-as 65001 neighbor FC00:172:16:255::4 update-source Loopback0 neighbor FC00:172:16:255::5 remote-as 65001 neighbor FC00:172:16:255::5 update-source Loopback0 ! 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 FC00:172:16:255::2 activate neighbor FC00:172:16:255::2 send-community both neighbor FC00:172:16:255::2 route-reflector-client neighbor FC00:172:16:255::4 activate neighbor FC00:172:16:255::4 send-community both neighbor FC00:172:16:255::4 route-reflector-client neighbor FC00:172:16:255::5 activate neighbor FC00:172:16:255::5 send-community both neighbor FC00:172:16:255::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 ! ipv6 router ospf 1 router-id 172.16.255.1 ! ipv6 pim rp-address FC00:172:16:255::255 ipv6 pim anycast-rp FC00:172:16:255::255 FC00:172:16:254::2 ! end </pre>	<pre> ! router ospf 1 router-id 172.16.255.2 ! router bgp 65001 bgp router-id 172.16.255.2 bgp log-neighbor-changes 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 FC00:172:16:255::1 remote-as 65001 neighbor FC00:172:16:255::1 update-source Loopback0 neighbor FC00:172:16:255::4 remote-as 65001 neighbor FC00:172:16:255::4 update-source Loopback0 neighbor FC00:172:16:255::5 remote-as 65001 neighbor FC00:172:16:255::5 update-source Loopback0 ! 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 FC00:172:16:255::1 activate neighbor FC00:172:16:255::1 send-community both neighbor FC00:172:16:255::1 route-reflector-client neighbor FC00:172:16:255::4 activate neighbor FC00:172:16:255::4 send-community both neighbor FC00:172:16:255::4 route-reflector-client neighbor FC00:172:16:255::5 activate neighbor FC00:172:16:255::5 send-community both neighbor FC00:172:16:255::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 ! ipv6 router ospf 1 router-id 172.16.255.2 ! ipv6 pim rp-address FC00:172:16:255::255 ipv6 pim anycast-rp FC00:172:16:255::255 FC00:172:16:254::1 ! end </pre>

Verifying BGP EVPN VXLAN with Dual Stack Underlay Configuration

The following sections provide sample output of **show** commands to verify the BGP EVPN VXLAN configuration with dual stack underlay.

- [Outputs to Verify Configuration on VTEP 1, on page 787](#)
- [Outputs to Verify Configuration on VTEP 2, on page 792](#)
- [Outputs to Verify Configuration on VTEP 3, on page 800](#)
- [Outputs to Verify Configuration on Spine 1, on page 807](#)
- [Outputs to Verify the Configuration on Spine 2, on page 813](#)

Outputs to Verify Configuration on VTEP 1

```
Leaf-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
       & - replicated local route overrides by connected

Gateway of last resort is not set

    172.16.0.0/16 is variably subnetted, 19 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, 02:27:03, GigabitEthernet1/0/1
O       172.16.15.0/24
        [110/2] via 172.16.13.1, 02:27:03, 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, 02:26:55, GigabitEthernet1/0/2
O       172.16.25.0/24
        [110/2] via 172.16.23.2, 02:26:55, GigabitEthernet1/0/2
O       172.16.254.1/32
        [110/2] via 172.16.13.1, 02:27:03, GigabitEthernet1/0/1
O       172.16.254.2/32
        [110/2] via 172.16.23.2, 02:26:55, 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, 02:26:55, GigabitEthernet1/0/2
        [110/3] via 172.16.13.1, 02:26:55, GigabitEthernet1/0/1
O       172.16.254.5/32
        [110/3] via 172.16.23.2, 02:26:55, GigabitEthernet1/0/2
        [110/3] via 172.16.13.1, 02:26:55, GigabitEthernet1/0/1
O       172.16.255.1/32
        [110/2] via 172.16.13.1, 02:27:03, GigabitEthernet1/0/1
O       172.16.255.2/32
        [110/2] via 172.16.23.2, 02:26:55, 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, 02:26:55, GigabitEthernet1/0/2
      [110/3] via 172.16.13.1, 02:26:55, GigabitEthernet1/0/1
O      172.16.255.5/32
      [110/3] via 172.16.23.2, 02:26:55, GigabitEthernet1/0/2
      [110/3] via 172.16.13.1, 02:26:55, GigabitEthernet1/0/1
O      172.16.255.255/32
      [110/2] via 172.16.23.2, 02:26:55, GigabitEthernet1/0/2
      [110/2] via 172.16.13.1, 02:27:03, GigabitEthernet1/0/1

```

Leaf-01# **show ipv6 route**

```

IPv6 Routing Table - default - 1 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
       I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
       EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
       NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
       OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
       ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
       ld - LISP dyn-eid, la - LISP away, le - LISP extranet-policy
       lp - LISP publications, ls - LISP destinations-summary
L  FF00::/8 [0/0]
   via Null0, receive

```

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, 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), 02:28:41/00:02:57, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Loopback1, Forward/Sparse, 02:28:39/00:02:23, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 02:26:54/00:02:50, flags:

(*, 225.0.0.101), 02:28:40/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, 02:28:40/stopped, flags:

(172.16.254.5, 225.0.0.101), 01:53:05/00:01:04, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 01:53:05/stopped, flags:

```

```
(172.16.254.4, 225.0.0.101), 02:20:24/00:01:01, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.23.2
  Outgoing interface list:
    Tunnel0, Forward/Sparse-Dense, 02:20:24/stopped, flags:

(172.16.254.3, 225.0.0.101), 02:28:09/00:03:22, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 02:25:53/00:02:32, flags: A
```

```
Leaf-01# show ipv6 mroute
No mroute entries found.
```

```
Leaf-01# show nve peer
```

```
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag
```

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/-/4	02:25:57
nve1	50901	L3CP	172.16.254.4	7c21.0dbd.9548	50901	UP	A/-/4	02:25:57
nve1	50901	L3CP	172.16.254.5	7c21.0dbd.2748	50901	UP	A/M/6	02:25:57
nve1	50901	L3CP	172.16.254.4	7c21.0dbd.9548	50901	UP	A/M/6	02:25:57
nve1	10101	L2CP	172.16.254.4	5	10101	UP	N/A	02:25:57
nve1	10101	L2CP	172.16.254.5	5	10101	UP	N/A	02:25:57
nve1	10102	L2CP	172.16.254.4	6	10102	UP	N/A	02:25:57
nve1	10102	L2CP	172.16.254.5	6	10102	UP	N/A	02:25:57

```
Leaf-01# 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
50 network entries using 19200 bytes of memory
68 path entries using 15776 bytes of memory
22/22 BGP path/bestpath attribute entries using 6512 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
16 BGP extended community entries using 864 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 42512 total bytes of memory
BGP activity 62/0 prefixes, 90/6 paths, scan interval 60 secs
50 networks peaked at 15:31:23 Oct 31 2022 UTC (00:00:34.727 ago)
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	186	174	54	0	0	02:26:54	18
172.16.255.2	4	65001	187	174	54	0	0	02:26:53	18

```
Leaf-01# 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][10B3D56A8FC1][32][10.1.101.1]/24	0.0.0.0			32768	?
*> [2][172.16.254.3:101][0][48][10B3D56A8FC1][128][FD00:10:1:101::1]/36					

```

0.0.0.0 32768 ?
*>i [2][172.16.254.3:101][0][48][44D3CA286CC1][0][*]/20
    172.16.254.4 0 100 0 ?
*>i [2][172.16.254.3:101][0][48][44D3CA286CC1][32][10.1.101.12]/24
    172.16.254.4 0 100 0 ?
*>i [2][172.16.254.3:101][0][48][44D3CA286CC3][0][*]/20
    172.16.254.5 0 100 0 ?
*>i [2][172.16.254.3:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
    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][7C210DBD2741][128][FD00:10:1:101::1]/36
    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 ?
*>i [2][172.16.254.3:101][0][48][7C210DBD9541][128][FD00:10:1:101::1]/36
    172.16.254.4 0 100 0 ?
*> [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
    0.0.0.0 32768 ?
*> [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
    0.0.0.0 32768 ?
*> [2][172.16.254.3:101][0][48][F4CFE24334C1][128][FE80::F6CF:E2FF:FE43:34C1]/36
    0.0.0.0 32768 ?
Route Distinguisher: 172.16.254.3:102
*> [2][172.16.254.3:102][0][48][10B3D56A8FCD][32][10.1.102.1]/24
    0.0.0.0 32768 ?
*> [2][172.16.254.3:102][0][48][10B3D56A8FCD][128][FD00:10:1:102::1]/36
    0.0.0.0 32768 ?
*>i [2][172.16.254.3:102][0][48][44D3CA286CC2][0][*]/20
    172.16.254.4 0 100 0 ?
*>i [2][172.16.254.3:102][0][48][44D3CA286CC2][32][10.1.102.12]/24
    172.16.254.4 0 100 0 ?
*>i [2][172.16.254.3:102][0][48][44D3CA286CC4][0][*]/20
    172.16.254.5 0 100 0 ?
*>i [2][172.16.254.3:102][0][48][44D3CA286CC4][32][10.1.102.13]/24
    172.16.254.5 0 100 0 ?
*>i [2][172.16.254.3:102][0][48][7C210DBD274D][32][10.1.102.1]/24
    172.16.254.5 0 100 0 ?
*>i [2][172.16.254.3:102][0][48][7C210DBD274D][128][FD00:10:1:102::1]/36
    172.16.254.5 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][FD00:10:1:102::1]/36
    172.16.254.4 0 100 0 ?
*> [2][172.16.254.3:102][0][48][F4CFE24334C2][0][*]/20
    0.0.0.0 32768 ?
*> [2][172.16.254.3:102][0][48][F4CFE24334C2][128][FE80::F6CF:E2FF:FE43:34C2]/36
    0.0.0.0 32768 ?
Route Distinguisher: 172.16.254.4:101
* i [2][172.16.254.4:101][0][48][44D3CA286CC1][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][44D3CA286CC1][32][10.1.101.12]/24
    172.16.254.4 0 100 0 ?
* i 172.16.254.4 0 100 0 ?
* 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 ?
* i [2][172.16.254.4:101][0][48][7C210DBD9541][128][FD00:10:1:101::1]/36
    172.16.254.4 0 100 0 ?
*>i 172.16.254.4 0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][0][*]/20
    172.16.254.4 0 100 0 ?

```

```

*>i          172.16.254.4          0    100    0 ?
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [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] [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] [FD00:10:1:102::1]/36
          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] [44D3CA286CC3] [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] [44D3CA286CC3] [32] [10.1.101.13]/24
          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 ?
* i [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
          172.16.254.5          0    100    0 ?
*>i          172.16.254.5          0    100    0 ?
Route Distinguisher: 172.16.254.5:102
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [0] [*]/20
          172.16.254.5          0    100    0 ?
*>i          172.16.254.5          0    100    0 ?
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
          172.16.254.5          0    100    0 ?
*>i          172.16.254.5          0    100    0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
          172.16.254.5          0    100    0 ?
*>i          172.16.254.5          0    100    0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
          172.16.254.5          0    100    0 ?
*>i          172.16.254.5          0    100    0 ?
Route Distinguisher: 172.16.254.3:102
*> [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
          0.0.0.0          0    100    32768 ?
*>i [3] [172.16.254.3:102] [0] [32] [172.16.254.4]/17
          172.16.254.4          0    100    0 ?
*>i [3] [172.16.254.3:102] [0] [32] [172.16.254.5]/17
          172.16.254.5          0    100    0 ?
Route Distinguisher: 172.16.254.4:102
* i [3] [172.16.254.4:102] [0] [32] [172.16.254.4]/17
          172.16.254.4          0    100    0 ?
*>i          172.16.254.4          0    100    0 ?
Route Distinguisher: 172.16.254.5:102
* i [3] [172.16.254.5:102] [0] [32] [172.16.254.5]/17
          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    100    32768 ?
*> [5] [1:1] [0] [24] [10.1.102.0]/17
          0.0.0.0          0    100    32768 ?
*> [5] [1:1] [0] [64] [FD00:10:1:101::]/29
          ::          0    100    32768 ?
*> [5] [1:1] [0] [64] [FD00:10:1:102::]/29
          ::          0    100    32768 ?

```

Leaf-01# show l2vpn evpn mac ip

IP Address	EVI	VLAN	MAC Address	Next Hop(s)
------------	-----	------	-------------	-------------

```

10.1.101.11          101  101  f4cf.e243.34c1 Gi1/0/4:101
10.1.101.12          101  101  44d3.ca28.6cc1 172.16.254.4
10.1.101.13          101  101  44d3.ca28.6cc3 172.16.254.5
FE80::F6CF:E2FF:FE43:34C1 101  101  f4cf.e243.34c1 Gi1/0/4:101
10.1.102.12          102  102  44d3.ca28.6cc2 172.16.254.4
10.1.102.13          102  102  44d3.ca28.6cc4 172.16.254.5
FE80::F6CF:E2FF:FE43:34C2 102  102  f4cf.e243.34c2 Gi1/0/5:102

```

```
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
          V1101:0
101      0  L2VPN 10b3.d56a.8fc1   FD00:10:1:101::1
          V1101:0
101      0  BGP 44d3.ca28.6cc1 10.1.101.12
          V:10101 172.16.254.4
101      0  BGP 44d3.ca28.6cc3 10.1.101.13
          V:10101 172.16.254.5
101      0  BGP 7c21.0dbd.2741 10.1.101.1
          V:10101 172.16.254.5
101      0  BGP 7c21.0dbd.2741  FD00:10:1:101::1
          V:10101 172.16.254.5
101      0  BGP 7c21.0dbd.9541 10.1.101.1
          V:10101 172.16.254.4
101      0  BGP 7c21.0dbd.9541  FD00:10:1:101::1
          V:10101 172.16.254.4
101      0  L2VPN f4cf.e243.34c1 10.1.101.11
          Gi1/0/4:101
101      0  L2VPN f4cf.e243.34c1  FE80::F6CF:E2FF:FE43:34C1
          Gi1/0/4:101
102      0  L2VPN 10b3.d56a.8fcd 10.1.102.1
          V1102:0
102      0  L2VPN 10b3.d56a.8fcd  FD00:10:1:102::1
          V1102:0
102      0  BGP 44d3.ca28.6cc2 10.1.102.12
          V:10102 172.16.254.4
102      0  BGP 44d3.ca28.6cc4 10.1.102.13
          V:10102 172.16.254.5
102      0  BGP 7c21.0dbd.274d 10.1.102.1
          V:10102 172.16.254.5
102      0  BGP 7c21.0dbd.274d  FD00:10:1:102::1
          V:10102 172.16.254.5
102      0  BGP 7c21.0dbd.954d 10.1.102.1
          V:10102 172.16.254.4
102      0  BGP 7c21.0dbd.954d  FD00:10:1:102::1
          V:10102 172.16.254.4
102      0  L2VPN f4cf.e243.34c2  FE80::F6CF:E2FF:FE43:34C2
          Gi1/0/5:102

```

To return to the example configuration, click [Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay, on page 777](#).

Outputs to Verify Configuration on VTEP 2

```
Leaf-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
 & - replicated local route overrides by connected

Gateway of last resort is not set

```

172.16.0.0/16 is variably subnetted, 19 subnets, 2 masks
O   172.16.13.0/24
    [110/2] via 172.16.14.1, 02:29:55, 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   172.16.15.0/24
    [110/2] via 172.16.14.1, 02:29:55, GigabitEthernet1/0/1
O   172.16.23.0/24
    [110/2] via 172.16.24.2, 02:29:52, 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   172.16.25.0/24
    [110/2] via 172.16.24.2, 02:29:52, GigabitEthernet1/0/2
O   172.16.254.1/32
    [110/2] via 172.16.14.1, 02:29:55, GigabitEthernet1/0/1
O   172.16.254.2/32
    [110/2] via 172.16.24.2, 02:29:52, GigabitEthernet1/0/2
O   172.16.254.3/32
    [110/3] via 172.16.24.2, 02:29:47, GigabitEthernet1/0/2
    [110/3] via 172.16.14.1, 02:29:55, GigabitEthernet1/0/1
C   172.16.254.4/32 is directly connected, Loopback1
O   172.16.254.5/32
    [110/3] via 172.16.24.2, 02:29:47, GigabitEthernet1/0/2
    [110/3] via 172.16.14.1, 02:29:47, GigabitEthernet1/0/1
O   172.16.255.1/32
    [110/2] via 172.16.14.1, 02:29:55, GigabitEthernet1/0/1
O   172.16.255.2/32
    [110/2] via 172.16.24.2, 02:29:52, GigabitEthernet1/0/2
O   172.16.255.3/32
    [110/3] via 172.16.24.2, 02:29:47, GigabitEthernet1/0/2
    [110/3] via 172.16.14.1, 02:29:55, GigabitEthernet1/0/1
C   172.16.255.4/32 is directly connected, Loopback0
O   172.16.255.5/32
    [110/3] via 172.16.24.2, 02:29:47, GigabitEthernet1/0/2
    [110/3] via 172.16.14.1, 02:29:47, GigabitEthernet1/0/1
O   172.16.255.255/32
    [110/2] via 172.16.24.2, 02:29:52, GigabitEthernet1/0/2
    [110/2] via 172.16.14.1, 02:29:55, GigabitEthernet1/0/1
  
```

Leaf-02# **show ipv6 route**

IPv6 Routing Table - default - 18 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
 B - BGP, R - RIP, H - NHRP, I1 - ISIS L1
 I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP
 EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination
 NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter
 OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1
 ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations
 ld - LISP dyn-eid, la - LISP away, le - LISP extranet-policy
 lp - LISP publications, ls - LISP destinations-summary

```

O FC00:172:16:13::/64 [110/2]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
C FC00:172:16:14::/64 [0/0]
  via GigabitEthernet1/0/1, directly connected
L FC00:172:16:14::4/128 [0/0]
  via GigabitEthernet1/0/1, receive
O FC00:172:16:15::/64 [110/2]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
O FC00:172:16:23::/64 [110/2]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
C FC00:172:16:24::/64 [0/0]
  via GigabitEthernet1/0/2, directly connected
L FC00:172:16:24::4/128 [0/0]
  via GigabitEthernet1/0/2, receive
O FC00:172:16:25::/64 [110/2]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
O FC00:172:16:254::1/128 [110/1]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
O FC00:172:16:254::2/128 [110/1]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
LC FC00:172:16:254::4/128 [0/0]
  via Loopback1, receive
O FC00:172:16:254::5/128 [110/2]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
O FC00:172:16:255::1/128 [110/1]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
O FC00:172:16:255::2/128 [110/1]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
LC FC00:172:16:255::4/128 [0/0]
  via Loopback0, receive
O FC00:172:16:255::5/128 [110/2]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
O FC00:172:16:255::255/128 [110/1]
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/1
  via FE80::7E21:DFF:FE92:B2D6, GigabitEthernet1/0/2
L FF00::/8 [0/0]
  via Null0, receive

```

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 - encaps-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), 02:30:33/00:02:28, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2

```



```

Outgoing interface list:
  Loopback1, Forward/Sparse, 02:30:32/00:02:28, flags:

(*, 225.0.0.101), 02:30:33/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnell1, Forward/Sparse-Dense, 02:30:33/stopped, flags:

(172.16.254.3, 225.0.0.101), 02:22:31/00:02:40, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnell1, Forward/Sparse-Dense, 02:22:31/stopped, flags:

(172.16.254.5, 225.0.0.101), 02:22:41/00:01:28, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.2
  Outgoing interface list:
    Tunnell1, Forward/Sparse-Dense, 02:22:41/stopped, flags:

(172.16.254.4, 225.0.0.101), 02:30:17/00:01:32, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0, Registering
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 02:29:51/00:03:28, flags: A

```

Leaf-02# **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

(*, FF0E::225::101), 02:30:33/never, RP FC00:172:16:255::255, flags: SCJ
  Incoming interface: GigabitEthernet1/0/2
  RPF nbr: FE80::7E21:DFF:FEBD:2CD6
  Immediate Outgoing interface list:
    Tunnel2, Forward, 02:30:33/never

```

Leaf-02# **show nve peer**

```

'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag

Interface VNI      Type Peer-IP      RMAC/Num_RT's  eVNI      state flags UP time
nve1     50901    L3CP 172.16.254.3  10b3.d56a.8fc8 50901     UP  A/-/4 02:28:51
nve1     50901    L3CP FC00:172:16:254::5 \
                                                7c21.0dbd.2748 50901     UP  A/-/4 02:28:51
nve1     50901    L3CP 172.16.254.3  10b3.d56a.8fc8 50901     UP  A/M/6 02:28:51
nve1     50901    L3CP FC00:172:16:254::5 \
                                                7c21.0dbd.2748 50901     UP  A/M/6 02:28:51
nve1     10101    L2CP 172.16.254.3   6           10101     UP  N/A   02:28:51
nve1     10101    L2CP FC00:172:16:254::5 \
                                                5           10101     UP  N/A   02:28:51
nve1     10102    L2CP 172.16.254.3   6           10102     UP  N/A   02:28:51
nve1     10102    L2CP FC00:172:16:254::5 \
                                                6           10102     UP  N/A   02:28:51

```

Leaf-02# show bgp l2vpn evpn summary

```

BGP router identifier 172.16.255.4, local AS number 65001
BGP table version is 76, main routing table version 76
51 network entries using 19584 bytes of memory
124 path entries using 28768 bytes of memory
23/22 BGP path/bestpath attribute entries using 6808 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
17 BGP extended community entries using 904 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 56224 total bytes of memory
BGP activity 64/0 prefixes, 188/43 paths, scan interval 60 secs
51 networks peaked at 15:31:23 Oct 31 2022 UTC (00:03:28.886 ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.1	4	65001	189	175	76	0	0	02:29:49	23
172.16.255.2	4	65001	189	175	76	0	0	02:29:45	23
FC00:172:16:255::1	4	65001	188	172	76	0	0	02:30:00	23
FC00:172:16:255::2	4	65001	189	173	76	0	0	02:29:58	23

Leaf-02# show bgp l2vpn evpn

```

BGP table version is 76, 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	172.16.254.3	0	100	0	?
* i [2]	[172.16.254.3:101]	[0]	[48]	[10B3D56A8FC1]	[128][FD00:10:1:101::1]/36
	172.16.254.3	0	100	0	?
* i	172.16.254.3	0	100	0	?
*>i	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	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	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	?
*>i	172.16.254.3	0	100	0	?
* i	172.16.254.3	0	100	0	?
Route Distinguisher: 172.16.254.3:102					
* i [2]	[172.16.254.3:102]	[0]	[48]	[10B3D56A8FCD]	[32][10.1.102.1]/24
	172.16.254.3	0	100	0	?
* i	172.16.254.3	0	100	0	?
*>i	172.16.254.3	0	100	0	?

```

* i          172.16.254.3          0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
          172.16.254.3          0 100 0 ?
* i          172.16.254.3          0 100 0 ?
*>i         172.16.254.3          0 100 0 ?
* i          172.16.254.3          0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20
          172.16.254.3          0 100 0 ?
* i          172.16.254.3          0 100 0 ?
*>i         172.16.254.3          0 100 0 ?
* i          172.16.254.3          0 100 0 ?
*>i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36
          172.16.254.3          0 100 0 ?
* i          172.16.254.3          0 100 0 ?
* i          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] [FD00:10:1:101::1]/36
          172.16.254.3          0 100 0 ?
*> [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [0] [*]/20
          0.0.0.0          32768 ?
*> [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [32] [10.1.101.12]/24
          0.0.0.0          32768 ?
*>i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC3] [0] [*]/20
          FC00:172:16:254::5
          0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC3] [32] [10.1.101.13]/24
          FC00:172:16:254::5
          0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [7C210DBD2741] [32] [10.1.101.1]/24
          FC00:172:16:254::5
          0 100 0 ?
*>i [2] [172.16.254.4:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
          FC00:172:16:254::5
          0 100 0 ?
*> [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
          0.0.0.0          32768 ?
*> [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
          0.0.0.0          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 [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] [10B3D56A8FCD] [32] [10.1.102.1]/24
          172.16.254.3          0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
          172.16.254.3          0 100 0 ?
*> [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [0] [*]/20
          0.0.0.0          32768 ?
*> [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
          0.0.0.0          32768 ?
*>i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC4] [0] [*]/20
          FC00:172:16:254::5
          0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
          FC00:172:16:254::5
          0 100 0 ?
*>i [2] [172.16.254.4:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
          FC00:172:16:254::5

```

```

0 100 0 ?
*>i [2][172.16.254.4:102][0][48][7C210DBD274D][128][FD00:10:1:102::1]/36
FC00:172:16:254::5
0 100 0 ?
*> [2][172.16.254.4:102][0][48][7C210DBD954D][32][10.1.102.1]/24
0.0.0.0 32768 ?
*> [2][172.16.254.4:102][0][48][7C210DBD954D][128][FD00:10:1:102::1]/36
0.0.0.0 32768 ?
*>i [2][172.16.254.4:102][0][48][F4CFE24334C2][0][*]/20
172.16.254.3 0 100 0 ?
*>i [2][172.16.254.4:102][0][48][F4CFE24334C2][128][FE80::F6CF:E2FF:FE43:34C2]/36
172.16.254.3 0 100 0 ?
Route Distinguisher: 172.16.254.5:101
* i [2][172.16.254.5:101][0][48][44D3CA286CC3][0][*]/20
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
* i [2][172.16.254.5:101][0][48][44D3CA286CC3][32][10.1.101.13]/24
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 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 ?
* i 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][128][FD00:10:1:101::1]/36
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [2][172.16.254.5:102][0][48][44D3CA286CC4][0][*]/20
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
* i [2][172.16.254.5:102][0][48][44D3CA286CC4][32][10.1.102.13]/24
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
* i [2][172.16.254.5:102][0][48][7C210DBD274D][32][10.1.102.1]/24
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
* i [2][172.16.254.5:102][0][48][7C210DBD274D][128][FD00:10:1:102::1]/36
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3][172.16.254.3:102][0][32][172.16.254.3]/17
172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
*>i 172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
Route Distinguisher: 172.16.254.4:102
*>i [3][172.16.254.4:102][0][32][172.16.254.3]/17
172.16.254.3 0 100 0 ?
*> [3][172.16.254.4:102][0][32][172.16.254.4]/17

```

```

0.0.0.0 32768 ?
*>i [3][172.16.254.4:102][0][32][172.16.254.5]/17
FC00:172:16:254::5 0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [3][172.16.254.5:102][0][32][172.16.254.5]/17
172.16.254.5 0 100 0 ?
* i 172.16.254.5 0 100 0 ?
* i 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 ?
* i 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][10.1.102.0]/17
172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
* i 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][FD00:10:1:101::]/29
172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
*> :: 0 32768 ?
* i [5][1:1][0][64][FD00:10:1:102::]/29
172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
* i 172.16.254.3 0 100 0 ?
*> :: 0 32768 ?

```

Leaf-02# show l2vpn evpn mac ip

IP Address	EVI	VLAN	MAC Address	Next Hop(s)
10.1.101.11	101	101	f4cf.e243.34c1	172.16.254.3
10.1.101.12	101	101	44d3.ca28.6cc1	Gi1/0/11:101
10.1.101.13	101	101	44d3.ca28.6cc3	FC00:172:16:254::5
FE80::F6CF:E2FF:FE43:34C1	101	101	f4cf.e243.34c1	172.16.254.3
10.1.102.12	102	102	44d3.ca28.6cc2	Gi1/0/12:102
10.1.102.13	102	102	44d3.ca28.6cc4	FC00:172:16:254::5
FE80::F6CF:E2FF:FE43:34C2	102	102	f4cf.e243.34c2	172.16.254.3

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	BGP	10b3.d56a.8fc1	FD00:10:1:101::1
		V:10101	172.16.254.3	
101	0	L2VPN	44d3.ca28.6cc1	10.1.101.12
			Gi1/0/11:101	
101	0	BGP	44d3.ca28.6cc3	10.1.101.13
		V:10101	FC00:172:16:254::5	
101	0	BGP	7c21.0dbd.2741	10.1.101.1
		V:10101	FC00:172:16:254::5	
101	0	BGP	7c21.0dbd.2741	FD00:10:1:101::1

```

V:10101 FC00:172:16:254::5
101      0 L2VPN 7c21.0dbd.9541          10.1.101.1
          Vl101:0
101      0 L2VPN 7c21.0dbd.9541          FD00:10:1:101::1
          Vl101:0
101      0 BGP f4cf.e243.34c1             10.1.101.11
          V:10101 172.16.254.3
101      0 BGP f4cf.e243.34c1             FE80::F6CF:E2FF:FE43:34C1
          V:10101 172.16.254.3
102      0 BGP 10b3.d56a.8fcd             10.1.102.1
          V:10102 172.16.254.3
102      0 BGP 10b3.d56a.8fcd             FD00:10:1:102::1
          V:10102 172.16.254.3
102      0 L2VPN 44d3.ca28.6cc2             10.1.102.12
          Gi1/0/12:102
102      0 BGP 44d3.ca28.6cc4             10.1.102.13
V:10102 FC00:172:16:254::5
102      0 BGP 7c21.0dbd.274d             10.1.102.1
V:10102 FC00:172:16:254::5
102      0 BGP 7c21.0dbd.274d             FD00:10:1:102::1
V:10102 FC00:172:16:254::5
102      0 L2VPN 7c21.0dbd.954d             10.1.102.1
          Vl102:0
102      0 L2VPN 7c21.0dbd.954d             FD00:10:1:102::1
          Vl102:0
102      0 BGP f4cf.e243.34c2             FE80::F6CF:E2FF:FE43:34C2
          V:10102 172.16.254.3

```

To return to the example configuration, click [Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay, on page 777](#).

Outputs to Verify Configuration on VTEP 3

Leaf-03# **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 Pfrp
       & - replicated local route overrides by connected

```

Gateway of last resort is not set

```

172.16.0.0/16 is variably subnetted, 19 subnets, 2 masks
O    172.16.13.0/24
     [110/2] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
O    172.16.14.0/24
     [110/2] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
C    172.16.15.0/24 is directly connected, GigabitEthernet1/0/1
L    172.16.15.5/32 is directly connected, GigabitEthernet1/0/1
O    172.16.23.0/24
     [110/2] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
O    172.16.24.0/24
     [110/2] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
C    172.16.25.0/24 is directly connected, GigabitEthernet1/0/2
L    172.16.25.5/32 is directly connected, GigabitEthernet1/0/2

```

```

O      172.16.254.1/32
      [110/2] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
O      172.16.254.2/32
      [110/2] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
O      172.16.254.3/32
      [110/3] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
      [110/3] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
O      172.16.254.4/32
      [110/3] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
      [110/3] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
C      172.16.254.5/32 is directly connected, Loopback1
O      172.16.255.1/32
      [110/2] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
O      172.16.255.2/32
      [110/2] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
O      172.16.255.3/32
      [110/3] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
      [110/3] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
O      172.16.255.4/32
      [110/3] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
      [110/3] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1
C      172.16.255.5/32 is directly connected, Loopback0
O      172.16.255.255/32
      [110/2] via 172.16.25.2, 02:30:22, GigabitEthernet1/0/2
      [110/2] via 172.16.15.1, 02:30:26, GigabitEthernet1/0/1

```

Leaf-03# **show ipv6 route**

IPv6 Routing Table - default - 18 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, R - RIP, H - NHRP, I1 - ISIS L1

I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP

EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination

Ndr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter

OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1

ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations

ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy

lp - LISP publications, ls - LISP destinations-summary

```

O      FC00:172:16:13::/64 [110/2]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
O      FC00:172:16:14::/64 [110/2]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
C      FC00:172:16:15::/64 [0/0]
      via GigabitEthernet1/0/1, directly connected
L      FC00:172:16:15::5/128 [0/0]
      via GigabitEthernet1/0/1, receive
O      FC00:172:16:23::/64 [110/2]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O      FC00:172:16:24::/64 [110/2]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
C      FC00:172:16:25::/64 [0/0]
      via GigabitEthernet1/0/2, directly connected
L      FC00:172:16:25::5/128 [0/0]
      via GigabitEthernet1/0/2, receive
O      FC00:172:16:254::1/128 [110/1]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
O      FC00:172:16:254::2/128 [110/1]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
O      FC00:172:16:254::4/128 [110/2]
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/1
      via FE80::7E21:DFE:FE92:B2D8, GigabitEthernet1/0/2
LC     FC00:172:16:254::5/128 [0/0]
      via Loopback1, receive
O      FC00:172:16:255::1/128 [110/1]

```

```

    via FE80::7E21:DFF:FE92:B2D8, GigabitEthernet1/0/1
O  FC00:172:16:255::2/128 [110/1]
    via FE80::7E21:DFF:FE92:B2D8, GigabitEthernet1/0/2
O  FC00:172:16:255::4/128 [110/2]
    via FE80::7E21:DFF:FE92:B2D8, GigabitEthernet1/0/1
    via FE80::7E21:DFF:FE92:B2D8, GigabitEthernet1/0/2
LC FC00:172:16:255::5/128 [0/0]
    via Loopback0, receive
O  FC00:172:16:255::255/128 [110/1]
    via FE80::7E21:DFF:FE92:B2D8, GigabitEthernet1/0/1
    via FE80::7E21:DFF:FE92:B2D8, GigabitEthernet1/0/2
L  FF00::/8 [0/0]
    via Null0, receive

```

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, 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), 02:30:48/00:02:16, RP 172.16.255.255, flags: SJCL
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.25.2
  Outgoing interface list:
    Loopback1, Forward/Sparse, 02:30:47/00:02:16, flags:

(*, 225.0.0.101), 02:30:48/stopped, RP 172.16.255.255, flags: SJCFx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.25.2
  Outgoing interface list:
    Tunnel1, Forward/Sparse-Dense, 02:30:48/stopped, flags:

(172.16.254.3, 225.0.0.101), 02:23:01/00:01:57, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.25.2
  Outgoing interface list:
    Tunnel1, Forward/Sparse-Dense, 02:23:01/stopped, flags:

(172.16.254.5, 225.0.0.101), 02:23:11/00:02:03, flags: FTx
  Incoming interface: Loopback1, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 02:23:11/00:02:56, flags: A

(172.16.254.4, 225.0.0.101), 02:23:49/00:00:40, flags: JTx
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.25.2
  Outgoing interface list:
    Tunnel1, Forward/Sparse-Dense, 02:23:49/stopped, flags:

```



```

Leaf-03# 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

(*, FF0E:225::101), 02:30:48/never, RP FC00:172:16:255::255, flags: SCJ
  Incoming interface: GigabitEthernet1/0/1
  RPF nbr: FE80::7E21:DFE:FE92:B2D8
  Immediate Outgoing interface list:
    Tunnel2, Forward, 02:30:48/never

Leaf-03# show nve peer
'M' - MAC entry download flag 'A' - Adjacency download flag
'4' - IPv4 flag '6' - IPv6 flag

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 02:29:22
nve1      50901    L3CP FC00:172:16:254::4 \
                                                7c21.0dbd.9548 50901      UP  A/-/4 02:29:22
nve1      50901    L3CP 172.16.254.3  10b3.d56a.8fc8 50901      UP  A/M/6 02:29:22
nve1      50901    L3CP FC00:172:16:254::4 \
                                                7c21.0dbd.9548 50901      UP  A/M/6 02:29:22
nve1      10101    L2CP 172.16.254.3  6            10101      UP  N/A   02:29:22
nve1      10101    L2CP FC00:172:16:254::4 \
                                                5            10101      UP  N/A   02:29:22
nve1      10102    L2CP 172.16.254.3  6            10102      UP  N/A   02:29:22
nve1      10102    L2CP FC00:172:16:254::4 \
                                                6            10102      UP  N/A   02:29:22

Leaf-03# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.5, local AS number 65001
BGP table version is 56, main routing table version 56
51 network entries using 19584 bytes of memory
78 path entries using 18096 bytes of memory
23/22 BGP path/bestpath attribute entries using 6808 bytes of memory
4 BGP rrinfo entries using 160 bytes of memory
17 BGP extended community entries using 904 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 45552 total bytes of memory
BGP activity 64/0 prefixes, 105/6 paths, scan interval 60 secs
51 networks peaked at 15:31:23 Oct 31 2022 UTC (00:03:59.747 ago)

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
FC00:172:16:255::1
4              65001    191    174     56     0   0 02:30:32    23
FC00:172:16:255::2
4              65001    191    174     56     0   0 02:30:25    23

Leaf-03# show bgp l2vpn evpn
BGP table version is 56, 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: 172.16.254.3:101
* 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][FD00:10:1:101::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][FE80::F6CF:E2FF:FE43:34C1]/36
    172.16.254.3          0 100 0 ?
*>i 172.16.254.3          0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [2][172.16.254.3:102][0][48][10B3D56A8FCD][32][10.1.102.1]/24
    172.16.254.3          0 100 0 ?
*>i 172.16.254.3          0 100 0 ?
* i [2][172.16.254.3:102][0][48][10B3D56A8FCD][128][FD00:10:1:102::1]/36
    172.16.254.3          0 100 0 ?
*>i 172.16.254.3          0 100 0 ?
* i [2][172.16.254.3:102][0][48][F4CFE24334C2][0][*]/20
    172.16.254.3          0 100 0 ?
*>i 172.16.254.3          0 100 0 ?
*>i [2][172.16.254.3:102][0][48][F4CFE24334C2][128][FE80::F6CF:E2FF:FE43:34C2]/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][44D3CA286CC1][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][44D3CA286CC1][32][10.1.101.12]/24
    172.16.254.4          0 100 0 ?
*>i 172.16.254.4          0 100 0 ?
* 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 ?
* i [2][172.16.254.4:101][0][48][7C210DBD9541][128][FD00:10:1:101::1]/36
    172.16.254.4          0 100 0 ?
*>i 172.16.254.4          0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][0][*]/20
    172.16.254.4          0 100 0 ?
*>i 172.16.254.4          0 100 0 ?
* i [2][172.16.254.4:102][0][48][44D3CA286CC2][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][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][FD00:10:1:102::1]/36
    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][10B3D56A8FC1][32][10.1.101.1]/24
    172.16.254.3          0 100 0 ?

```

```

*>i [2] [172.16.254.5:101] [0] [48] [10B3D56A8FC1] [128] [FD00:10:1:101::1]/36
      172.16.254.3          0      100      0 ?
*>i [2] [172.16.254.5:101] [0] [48] [44D3CA286CC1] [0] [*]/20
      FC00:172:16:254::4
                                     0      100      0 ?
*>i [2] [172.16.254.5:101] [0] [48] [44D3CA286CC1] [32] [10.1.101.12]/24
      FC00:172:16:254::4
                                     0      100      0 ?
*> [2] [172.16.254.5:101] [0] [48] [44D3CA286CC3] [0] [*]/20
      0.0.0.0                                32768 ?
*> [2] [172.16.254.5:101] [0] [48] [44D3CA286CC3] [32] [10.1.101.13]/24
      0.0.0.0                                32768 ?
*> [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [32] [10.1.101.1]/24
      0.0.0.0                                32768 ?
*> [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
      0.0.0.0                                32768 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD9541] [32] [10.1.101.1]/24
      FC00:172:16:254::4
                                     0      100      0 ?
*>i [2] [172.16.254.5:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
      FC00:172:16:254::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 [2] [172.16.254.5:101] [0] [48] [F4CFE24334C1] [128] [FE80::F6CF:E2FF:FE43:34C1]/36
      172.16.254.3          0      100      0 ?
Route Distinguisher: 172.16.254.5:102
*>i [2] [172.16.254.5:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24
      172.16.254.3          0      100      0 ?
*>i [2] [172.16.254.5:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
      172.16.254.3          0      100      0 ?
*>i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC2] [0] [*]/20
      FC00:172:16:254::4
                                     0      100      0 ?
*>i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
      FC00:172:16:254::4
                                     0      100      0 ?
*> [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [0] [*]/20
      0.0.0.0                                32768 ?
*> [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
      0.0.0.0                                32768 ?
*> [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24
      0.0.0.0                                32768 ?
*> [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1]/36
      0.0.0.0                                32768 ?
*>i [2] [172.16.254.5:102] [0] [48] [7C210DBD954D] [32] [10.1.102.1]/24
      FC00:172:16:254::4
                                     0      100      0 ?
*>i [2] [172.16.254.5:102] [0] [48] [7C210DBD954D] [128] [FD00:10:1:102::1]/36
      FC00:172:16:254::4
                                     0      100      0 ?
*>i [2] [172.16.254.5:102] [0] [48] [F4CFE24334C2] [0] [*]/20
      172.16.254.3          0      100      0 ?
*>i [2] [172.16.254.5:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36
      172.16.254.3          0      100      0 ?
Route Distinguisher: 172.16.254.3:102
* i [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
      172.16.254.3          0      100      0 ?
*>i [3] [172.16.254.3:102] [0] [32] [172.16.254.3]/17
      172.16.254.3          0      100      0 ?
Route Distinguisher: 172.16.254.4:102
* i [3] [172.16.254.4:102] [0] [32] [172.16.254.4]/17
      172.16.254.4          0      100      0 ?

```

```

*>i          172.16.254.4          0    100    0 ?
Route Distinguisher: 172.16.254.5:102
*>i [3][172.16.254.5:102][0][32][172.16.254.3]/17
    172.16.254.3          0    100    0 ?
*>i [3][172.16.254.5:102][0][32][172.16.254.4]/17
    FC00:172:16:254::4          0    100    0 ?
*> [3][172.16.254.5:102][0][32][172.16.254.5]/17
    0.0.0.0          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 ?
*>          0.0.0.0          0    32768 ?
* i [5][1:1][0][24][10.1.102.0]/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][64][FD00:10:1:101::]/29
    172.16.254.3          0    100    0 ?
* i          172.16.254.3          0    100    0 ?
*>          ::          0    32768 ?
* i [5][1:1][0][64][FD00:10:1:102::]/29
    172.16.254.3          0    100    0 ?
* i          172.16.254.3          0    100    0 ?
*>          ::          0    32768 ?

```

Leaf-03# show l2vpn evpn mac ip

IP Address	EVI	VLAN	MAC Address	Next Hop(s)
10.1.101.11	101	101	f4cf.e243.34c1	172.16.254.3
10.1.101.12	101	101	44d3.ca28.6cc1	FC00:172:16:254::4
10.1.101.13	101	101	44d3.ca28.6cc3	Gi1/0/16:101
FE80::F6CF:E2FF:FE43:34C1	101	101	f4cf.e243.34c1	172.16.254.3
10.1.102.12	102	102	44d3.ca28.6cc2	FC00:172:16:254::4
10.1.102.13	102	102	44d3.ca28.6cc4	Gi1/0/17:102
FE80::F6CF:E2FF:FE43:34C2	102	102	f4cf.e243.34c2	172.16.254.3

Leaf-03# show l2route evpn mac ip

EVI	Etag	Prod	Mac Address	Next Hop(s)	Host IP
101	0	BGP	10b3.d56a.8fc1		10.1.101.1
			V:10101 172.16.254.3		
101	0	BGP	10b3.d56a.8fc1		FD00:10:1:101::1
			V:10101 172.16.254.3		
101	0	BGP	44d3.ca28.6cc1		10.1.101.12
			V:10101 FC00:172:16:254::4		
101	0	L2VPN	44d3.ca28.6cc3		10.1.101.13
			Gi1/0/16:101		
101	0	L2VPN	7c21.0dbd.2741		10.1.101.1
			Vl101:0		
101	0	L2VPN	7c21.0dbd.2741		FD00:10:1:101::1
			Vl101:0		
101	0	BGP	7c21.0dbd.9541		10.1.101.1
			V:10101 FC00:172:16:254::4		
101	0	BGP	7c21.0dbd.9541		FD00:10:1:101::1
			V:10101 FC00:172:16:254::4		
101	0	BGP	f4cf.e243.34c1		10.1.101.11
			V:10101 172.16.254.3		

```

101      0   BGP f4cf.e243.34c1                FE80::F6CF:E2FF:FE43:34C1
          V:10101 172.16.254.3
102      0   BGP 10b3.d56a.8fcd                10.1.1.102.1
          V:10102 172.16.254.3
102      0   BGP 10b3.d56a.8fcd                FD00:10:1:102::1
          V:10102 172.16.254.3
102      0   BGP 44d3.ca28.6cc2                10.1.1.102.12
          V:10102 FC00:172:16:254::4
102      0   L2VPN 44d3.ca28.6cc4              10.1.1.102.13
          Gi1/0/17:102
102      0   L2VPN 7c21.0dbd.274d              10.1.1.102.1
          V1102:0
102      0   L2VPN 7c21.0dbd.274d              FD00:10:1:102::1
          V1102:0
102      0   BGP 7c21.0dbd.954d                10.1.1.102.1
          V:10102 FC00:172:16:254::4
102      0   BGP 7c21.0dbd.954d                FD00:10:1:102::1
          V:10102 FC00:172:16:254::4
102      0   BGP f4cf.e243.34c2                FE80::F6CF:E2FF:FE43:34C2
          V:10102 172.16.254.3

```

To return to the example configuration, click [Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay, on page 777](#).

Outputs to Verify Configuration on Spine 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
       & - replicated local route overrides by connected

Gateway of last resort is not set

172.16.0.0/16 is variably subnetted, 20 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, 03:11:49, GigabitEthernet1/0/1
O       172.16.24.0/24
         [110/2] via 172.16.14.4, 03:11:49, GigabitEthernet1/0/2
O       172.16.25.0/24
         [110/2] via 172.16.15.5, 03:11:48, GigabitEthernet1/0/3
C       172.16.254.1/32 is directly connected, Loopback1
O       172.16.254.2/32
         [110/3] via 172.16.15.5, 03:11:39, GigabitEthernet1/0/3
         [110/3] via 172.16.14.4, 03:11:43, GigabitEthernet1/0/2
         [110/3] via 172.16.13.3, 03:11:43, GigabitEthernet1/0/1
O       172.16.254.3/32
         [110/2] via 172.16.13.3, 03:11:49, GigabitEthernet1/0/1
O       172.16.254.4/32

```

```

O      [110/2] via 172.16.14.4, 03:11:49, GigabitEthernet1/0/2
172.16.254.5/32
C      [110/2] via 172.16.15.5, 03:11:48, GigabitEthernet1/0/3
172.16.255.1/32 is directly connected, Loopback0
O      172.16.255.2/32
      [110/3] via 172.16.15.5, 03:11:39, GigabitEthernet1/0/3
      [110/3] via 172.16.14.4, 03:11:43, GigabitEthernet1/0/2
      [110/3] via 172.16.13.3, 03:11:43, GigabitEthernet1/0/1
O      172.16.255.3/32
      [110/2] via 172.16.13.3, 03:11:49, GigabitEthernet1/0/1
O      172.16.255.4/32
      [110/2] via 172.16.14.4, 03:11:49, GigabitEthernet1/0/2
O      172.16.255.5/32
      [110/2] via 172.16.15.5, 03:11:48, GigabitEthernet1/0/3
C      172.16.255.255/32 is directly connected, Loopback2

```

Spine-01# **show ipv6 route**

IPv6 Routing Table - default - 19 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, R - RIP, H - NHRP, I1 - ISIS L1

I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP

EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination

NDr - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter

OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1

ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations

ld - LISP dyn-eid, lA - LISP away, le - LISP extranet-policy

lp - LISP publications, ls - LISP destinations-summary

```

C FC00:172:16:13::/64 [0/0]
  via GigabitEthernet1/0/1, directly connected
L FC00:172:16:13::1/128 [0/0]
  via GigabitEthernet1/0/1, receive
C FC00:172:16:14::/64 [0/0]
  via GigabitEthernet1/0/2, directly connected
L FC00:172:16:14::1/128 [0/0]
  via GigabitEthernet1/0/2, receive
C FC00:172:16:15::/64 [0/0]
  via GigabitEthernet1/0/3, directly connected
L FC00:172:16:15::1/128 [0/0]
  via GigabitEthernet1/0/3, receive
O FC00:172:16:23::/64 [110/3]
  via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
  via FE80::7E21:DFF:FEBD:2764, GigabitEthernet1/0/3
O FC00:172:16:24::/64 [110/2]
  via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
O FC00:172:16:25::/64 [110/2]
  via FE80::7E21:DFF:FEBD:2764, GigabitEthernet1/0/3
LC FC00:172:16:254::1/128 [0/0]
  via Loopback1, receive
O FC00:172:16:254::2/128 [110/2]
  via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
  via FE80::7E21:DFF:FEBD:2764, GigabitEthernet1/0/3
O FC00:172:16:254::4/128 [110/1]
  via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
O FC00:172:16:254::5/128 [110/1]
  via FE80::7E21:DFF:FEBD:2764, GigabitEthernet1/0/3
LC FC00:172:16:255::1/128 [0/0]
  via Loopback0, receive
O FC00:172:16:255::2/128 [110/2]
  via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
  via FE80::7E21:DFF:FEBD:2764, GigabitEthernet1/0/3
O FC00:172:16:255::4/128 [110/1]
  via FE80::7E21:DFF:FEBD:9564, GigabitEthernet1/0/2
O FC00:172:16:255::5/128 [110/1]

```

```

    via FE80::7E21:DFE:FEBD:2764, GigabitEthernet1/0/3
LC FC00:172:16:255::255/128 [0/0]
    via Loopback2, receive
L FF00::/8 [0/0]
    via Null0, receive

```

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), 03:11:59/00:02:07, RP 172.16.255.255, flags: SPL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list: Null

(*, 225.0.0.101), 00:19:24/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), 00:01:27/00:01:32, flags: PA
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.14.4
  Outgoing interface list: Null

(172.16.254.5, 225.0.0.101), 00:01:56/00:01:03, flags: PA
  Incoming interface: GigabitEthernet1/0/3, RPF nbr 172.16.15.5
  Outgoing interface list: Null

```

Spine-01# **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

```

(*, FF0E:225::101), 03:11:56/00:02:40, RP FC00:172:16:255::255, flags: S
  Incoming interface: Tunnel2
  RPF nbr: FC00:172:16:255::255
  Immediate Outgoing interface list:

```

GigabitEthernet1/0/3, Forward, 03:11:54/00:02:40

Spine-01# **show bgp l2vpn evpn summary**

BGP router identifier 172.16.255.1, local AS number 65001
 BGP table version is 33, main routing table version 33
 32 network entries using 12288 bytes of memory
 85 path entries using 19720 bytes of memory
 24/22 BGP path/bestpath attribute entries using 7104 bytes of memory
 3 BGP rrinfo entries using 120 bytes of memory
 17 BGP extended community entries using 920 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 40152 total bytes of memory
 BGP activity 32/0 prefixes, 85/0 paths, scan interval 60 secs
 32 networks peaked at 15:31:23 Oct 31 2022 UTC (00:45:20.867 ago)

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.3	4	65001	224	236	33	0	0	03:11:40	14
172.16.255.4	4	65001	222	234	33	0	0	03:11:41	13
FC00:172:16:255::2	4	65001	233	235	33	0	0	03:11:50	32
FC00:172:16:255::4	4	65001	219	234	33	0	0	03:11:52	13
FC00:172:16:255::5	4	65001	220	238	33	0	0	03:11:53	13

Spine-01# **show bgp l2vpn evpn**

BGP table version is 33, 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] [FD00:10:1:101::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] [FE80::F6CF:E2FF:FE43:34C1]/36	172.16.254.3	0	100	0 ?	
*>i	172.16.254.3	0	100	0 ?	
Route Distinguisher: 172.16.254.3:102					
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24	172.16.254.3	0	100	0 ?	
*>i	172.16.254.3	0	100	0 ?	
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36	172.16.254.3	0	100	0 ?	
*>i	172.16.254.3	0	100	0 ?	
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20	172.16.254.3	0	100	0 ?	
*>i	172.16.254.3	0	100	0 ?	
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/36	172.16.254.3	0	100	0 ?	


```

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
172.16.254.4          0 100 0 ?
* i 172.16.254.4      0 100 0 ?
*>i 172.16.254.4      0 100 0 ?
* i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [32] [10.1.101.12] /24
172.16.254.4          0 100 0 ?
* i 172.16.254.4      0 100 0 ?
*>i 172.16.254.4      0 100 0 ?
* 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 ?
*>i 172.16.254.4      0 100 0 ?
* i [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1] /36
172.16.254.4          0 100 0 ?
* i 172.16.254.4      0 100 0 ?
*>i 172.16.254.4      0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [0] [*] /20
172.16.254.4          0 100 0 ?
* i 172.16.254.4      0 100 0 ?
*>i 172.16.254.4      0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12] /24
172.16.254.4          0 100 0 ?
* i 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 172.16.254.4      0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [128] [FD00:10:1:102::1] /36
172.16.254.4          0 100 0 ?
* i 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] [44D3CA286CC3] [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] [44D3CA286CC3] [32] [10.1.101.13] /24
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 ?
* i [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1] /36
172.16.254.5          0 100 0 ?
*>i 172.16.254.5      0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [0] [*] /20
172.16.254.5          0 100 0 ?
*>i 172.16.254.5      0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13] /24
172.16.254.5          0 100 0 ?
*>i 172.16.254.5      0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1] /24
172.16.254.5          0 100 0 ?
*>i 172.16.254.5      0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [128] [FD00:10:1:102::1] /36
172.16.254.5          0 100 0 ?
*>i 172.16.254.5      0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3] [172.16.254.3:102] [0] [32] [172.16.254.3] /17

```

```

172.16.254.3          0 100 0 ?
*>i 172.16.254.3      0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [3][172.16.254.4:102][0][32][172.16.254.4]/17
172.16.254.4        0 100 0 ?
* i 172.16.254.4     0 100 0 ?
*>i 172.16.254.4     0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [3][172.16.254.5:102][0][32][172.16.254.5]/17
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
172.16.254.3        0 100 0 ?
* i 172.16.254.5     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
*>i 172.16.254.3     0 100 0 ?
* i [5][1:1][0][24][10.1.102.0]/17
172.16.254.3        0 100 0 ?
* i 172.16.254.5     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
*>i 172.16.254.3     0 100 0 ?
* i [5][1:1][0][64][FD00:10:1:101::]/29
172.16.254.3        0 100 0 ?
* i 172.16.254.5     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
*>i 172.16.254.3     0 100 0 ?
* i [5][1:1][0][64][FD00:10:1:102::]/29
172.16.254.3        0 100 0 ?
* i 172.16.254.5     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
* i 172.16.254.4     0 100 0 ?
*>i 172.16.254.3     0 100 0 ?

```

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): 03:11:04, Messages sent/received: 249/260
Output messages discarded: 0
Connection and counters cleared 03:12:04 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: 3
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: 249/156
SA Requests in: 0
SA Responses out: 0
Data Packets in/out: 36/43

```
Spine-01# show ip msdp sa-cache
MSDP Source-Active Cache - 3 entries
(172.16.254.3, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 03:07:36/00:05:34, Peer 172.16.254.2
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:04:37/00:03:42, Peer 172.16.254.2
(172.16.254.5, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:50:54/00:00:50, Peer 172.16.254.2
```

```
Spine-01# show ipv6 pim anycast-rp
Anycast RP Peers For FC00:172:16:255::255 Last Register/Register-Stop received
FC00:172:16:254::2 03:12:04/03:12:04
```

To return to the example configuration, click [Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay, on page 777](#).

Outputs to Verify the Configuration on Spine 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
& - replicated local route overrides by connected
```

Gateway of last resort is not set

```
172.16.0.0/16 is variably subnetted, 20 subnets, 2 masks
O    172.16.13.0/24
     [110/2] via 172.16.23.3, 03:12:16, GigabitEthernet1/0/1
O    172.16.14.0/24
     [110/2] via 172.16.24.4, 03:12:17, GigabitEthernet1/0/2
O    172.16.15.0/24
     [110/2] via 172.16.25.5, 03:12:16, 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    172.16.254.1/32
     [110/3] via 172.16.25.5, 03:12:16, GigabitEthernet1/0/3
     [110/3] via 172.16.24.4, 03:12:17, GigabitEthernet1/0/2
     [110/3] via 172.16.23.3, 03:12:16, GigabitEthernet1/0/1
C    172.16.254.2/32 is directly connected, Loopback1
O    172.16.254.3/32
     [110/2] via 172.16.23.3, 03:12:16, GigabitEthernet1/0/1
O    172.16.254.4/32
     [110/2] via 172.16.24.4, 03:12:17, GigabitEthernet1/0/2
O    172.16.254.5/32
     [110/2] via 172.16.25.5, 03:12:16, GigabitEthernet1/0/3
O    172.16.255.1/32
     [110/3] via 172.16.25.5, 03:12:16, GigabitEthernet1/0/3
```

```

    [110/3] via 172.16.24.4, 03:12:17, GigabitEthernet1/0/2
    [110/3] via 172.16.23.3, 03:12:16, 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, 03:12:16, GigabitEthernet1/0/1
O   172.16.255.4/32
    [110/2] via 172.16.24.4, 03:12:17, GigabitEthernet1/0/2
O   172.16.255.5/32
    [110/2] via 172.16.25.5, 03:12:16, GigabitEthernet1/0/3
C   172.16.255.255/32 is directly connected, Loopback2

```

Spine-02# **show ipv6 route**

IPv6 Routing Table - default - 19 entries

Codes: C - Connected, L - Local, S - Static, U - Per-user Static route

B - BGP, R - RIP, H - NHRP, I1 - ISIS L1

I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary, D - EIGRP

EX - EIGRP external, ND - ND Default, NDp - ND Prefix, DCE - Destination

NDR - Redirect, RL - RPL, O - OSPF Intra, OI - OSPF Inter

OE1 - OSPF ext 1, OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1

ON2 - OSPF NSSA ext 2, la - LISP alt, lr - LISP site-registrations

ld - LISP dyn-eid, la - LISP away, le - LISP extranet-policy

lp - LISP publications, ls - LISP destinations-summary

```

O   FC00:172:16:13::/64 [110/3]
    via FE80::7E21:DFE:FEED:9556, GigabitEthernet1/0/2
    via FE80::7E21:DFE:FEED:2756, GigabitEthernet1/0/3
O   FC00:172:16:14::/64 [110/2]
    via FE80::7E21:DFE:FEED:9556, GigabitEthernet1/0/2
O   FC00:172:16:15::/64 [110/2]
    via FE80::7E21:DFE:FEED:2756, GigabitEthernet1/0/3
C   FC00:172:16:23::/64 [0/0]
    via GigabitEthernet1/0/1, directly connected
L   FC00:172:16:23::2/128 [0/0]
    via GigabitEthernet1/0/1, receive
C   FC00:172:16:24::/64 [0/0]
    via GigabitEthernet1/0/2, directly connected
L   FC00:172:16:24::2/128 [0/0]
    via GigabitEthernet1/0/2, receive
C   FC00:172:16:25::/64 [0/0]
    via GigabitEthernet1/0/3, directly connected
L   FC00:172:16:25::2/128 [0/0]
    via GigabitEthernet1/0/3, receive
O   FC00:172:16:254::1/128 [110/2]
    via FE80::7E21:DFE:FEED:9556, GigabitEthernet1/0/2
    via FE80::7E21:DFE:FEED:2756, GigabitEthernet1/0/3
LC  FC00:172:16:254::2/128 [0/0]
    via Loopback1, receive
O   FC00:172:16:254::4/128 [110/1]
    via FE80::7E21:DFE:FEED:9556, GigabitEthernet1/0/2
O   FC00:172:16:254::5/128 [110/1]
    via FE80::7E21:DFE:FEED:2756, GigabitEthernet1/0/3
O   FC00:172:16:255::1/128 [110/2]
    via FE80::7E21:DFE:FEED:9556, GigabitEthernet1/0/2
    via FE80::7E21:DFE:FEED:2756, GigabitEthernet1/0/3
LC  FC00:172:16:255::2/128 [0/0]
    via Loopback0, receive
O   FC00:172:16:255::4/128 [110/1]
    via FE80::7E21:DFE:FEED:9556, GigabitEthernet1/0/2
O   FC00:172:16:255::5/128 [110/1]
    via FE80::7E21:DFE:FEED:2756, GigabitEthernet1/0/3
LC  FC00:172:16:255::255/128 [0/0]
    via Loopback2, receive
L   FF00::/8 [0/0]
    via Null0, receive

```

```

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), 03:12:27/00:03:29, RP 172.16.255.255, flags: SJCL
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/3, Forward/Sparse, 03:12:16/00:03:29, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 03:12:16/00:03:17, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 03:12:27/00:03:14, flags:

(*, 225.0.0.101), 03:12:26/00:03:19, RP 172.16.255.255, flags: S
  Incoming interface: Null, RPF nbr 0.0.0.0
  Outgoing interface list:
    GigabitEthernet1/0/1, Forward/Sparse, 03:12:13/00:03:19, flags:
    GigabitEthernet1/0/3, Forward/Sparse, 03:12:16/00:02:39, flags:
    GigabitEthernet1/0/2, Forward/Sparse, 03:12:16/00:02:31, flags:

(172.16.254.3, 225.0.0.101), 03:05:05/00:02:26, flags: TA
  Incoming interface: GigabitEthernet1/0/1, RPF nbr 172.16.23.3
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 03:05:05/00:02:37, flags:
    GigabitEthernet1/0/3, Forward/Sparse, 03:05:05/00:02:39, flags:

(172.16.254.5, 225.0.0.101), 03:05:06/00:03:09, flags: T
  Incoming interface: GigabitEthernet1/0/3, RPF nbr 172.16.25.5
  Outgoing interface list:
    GigabitEthernet1/0/2, Forward/Sparse, 03:05:06/00:03:23, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 03:05:06/00:03:19, flags:

(172.16.254.4, 225.0.0.101), 03:09:13/00:00:15, flags: T
  Incoming interface: GigabitEthernet1/0/2, RPF nbr 172.16.24.4
  Outgoing interface list:
    GigabitEthernet1/0/3, Forward/Sparse, 03:09:13/00:02:43, flags:
    GigabitEthernet1/0/1, Forward/Sparse, 03:09:13/00:03:19, flags:

```

```

Spine-02# 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

(*, FF0E:225::101), 03:12:23/00:03:11, RP FC00:172:16:255::255, flags: S
  Incoming interface: Tunnel2
  RPF nbr: FC00:172:16:255::255
  Immediate Outgoing interface list:
    GigabitEthernet1/0/2, Forward, 03:12:23/00:03:11

```

```

Spine-02# show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 33, main routing table version 33
32 network entries using 12288 bytes of memory
85 path entries using 19720 bytes of memory
24/22 BGP path/bestpath attribute entries using 7104 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
17 BGP extended community entries using 920 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 40152 total bytes of memory
BGP activity 32/0 prefixes, 85/0 paths, scan interval 60 secs
32 networks peaked at 15:31:23 Oct 31 2022 UTC (00:45:53.775 ago)

```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
172.16.255.3	4	65001	224	237	33	0	0	03:12:12	14
172.16.255.4	4	65001	223	235	33	0	0	03:12:10	13
FC00:172:16:255::1									
	4	65001	235	234	33	0	0	03:12:22	32
FC00:172:16:255::4									
	4	65001	220	236	33	0	0	03:12:22	13
FC00:172:16:255::5									
	4	65001	220	237	33	0	0	03:12:19	13

```

Spine-02# show bgp l2vpn evpn
BGP table version is 33, 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] [FD00:10:1:101::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] [FE80::F6CF:E2FF:FE43:34C1]/36	172.16.254.3	0	100	0	?
*>i	172.16.254.3	0	100	0	?
Route Distinguisher: 172.16.254.3:102					
* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [32] [10.1.102.1]/24	172.16.254.3	0	100	0	?
*>i	172.16.254.3	0	100	0	?

```

* i [2] [172.16.254.3:102] [0] [48] [10B3D56A8FCD] [128] [FD00:10:1:102::1]/36
    172.16.254.3 0 100 0 ?
*>i 172.16.254.3 0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [0] [*]/20
    172.16.254.3 0 100 0 ?
*>i 172.16.254.3 0 100 0 ?
* i [2] [172.16.254.3:102] [0] [48] [F4CFE24334C2] [128] [FE80::F6CF:E2FF:FE43:34C2]/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] [44D3CA286CC1] [0] [*]/20
    172.16.254.4 0 100 0 ?
* i 172.16.254.4 0 100 0 ?
*>i 172.16.254.4 0 100 0 ?
* i [2] [172.16.254.4:101] [0] [48] [44D3CA286CC1] [32] [10.1.101.12]/24
    172.16.254.4 0 100 0 ?
* i 172.16.254.4 0 100 0 ?
*>i 172.16.254.4 0 100 0 ?
* 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 ?
*>i 172.16.254.4 0 100 0 ?
* i [2] [172.16.254.4:101] [0] [48] [7C210DBD9541] [128] [FD00:10:1:101::1]/36
    172.16.254.4 0 100 0 ?
* i 172.16.254.4 0 100 0 ?
*>i 172.16.254.4 0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [0] [*]/20
    172.16.254.4 0 100 0 ?
* i 172.16.254.4 0 100 0 ?
*>i 172.16.254.4 0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [44D3CA286CC2] [32] [10.1.102.12]/24
    172.16.254.4 0 100 0 ?
* i 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 172.16.254.4 0 100 0 ?
* i [2] [172.16.254.4:102] [0] [48] [7C210DBD954D] [128] [FD00:10:1:102::1]/36
    172.16.254.4 0 100 0 ?
* i 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] [44D3CA286CC3] [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] [44D3CA286CC3] [32] [10.1.101.13]/24
    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 ?
* i [2] [172.16.254.5:101] [0] [48] [7C210DBD2741] [128] [FD00:10:1:101::1]/36
    172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [0] [*]/20
    172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [44D3CA286CC4] [32] [10.1.102.13]/24
    172.16.254.5 0 100 0 ?
*>i 172.16.254.5 0 100 0 ?
* i [2] [172.16.254.5:102] [0] [48] [7C210DBD274D] [32] [10.1.102.1]/24

```

```

172.16.254.5          0 100 0 ?
*>i          172.16.254.5          0 100 0 ?
* i [2][172.16.254.5:102][0][48][7C210DBD274D][128][FD00:10:1:102::1]/36
172.16.254.5          0 100 0 ?
*>i          172.16.254.5          0 100 0 ?
Route Distinguisher: 172.16.254.3:102
* i [3][172.16.254.3:102][0][32][172.16.254.3]/17
172.16.254.3          0 100 0 ?
*>i          172.16.254.3          0 100 0 ?
Route Distinguisher: 172.16.254.4:102
* i [3][172.16.254.4:102][0][32][172.16.254.4]/17
172.16.254.4          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
*>i          172.16.254.4          0 100 0 ?
Route Distinguisher: 172.16.254.5:102
* i [3][172.16.254.5:102][0][32][172.16.254.5]/17
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
172.16.254.3          0 100 0 ?
* i          172.16.254.5          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
*>i          172.16.254.3          0 100 0 ?
* i [5][1:1][0][24][10.1.102.0]/17
172.16.254.3          0 100 0 ?
* i          172.16.254.5          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
*>i          172.16.254.3          0 100 0 ?
* i [5][1:1][0][64][FD00:10:1:101::]/29
172.16.254.3          0 100 0 ?
* i          172.16.254.5          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
*>i          172.16.254.3          0 100 0 ?
* i [5][1:1][0][64][FD00:10:1:102::]/29
172.16.254.3          0 100 0 ?
* i          172.16.254.5          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
* i          172.16.254.4          0 100 0 ?
*>i          172.16.254.3          0 100 0 ?

```

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): 03:11:40, Messages sent/received: 261/250

Output messages discarded: 0

Connection and counters cleared 03:12:35 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: 200/196
SA Requests in: 0
SA Responses out: 0
Data Packets in/out: 37/54
```

```
Spine-02# show ip msdp sa-cache
MSDP Source-Active Cache - 2 entries
(172.16.254.4, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 01:18:26/00:05:40, Peer 172.16.254.1
(172.16.254.5, 225.0.0.101), RP 172.16.255.255, BGP/AS 0, 00:02:33/00:05:40, Peer 172.16.254.1
```

```
Spine-02# show ipv6 pim anycast-rp
Anycast RP Peers For FC00:172:16:255::255 Last Register/Register-Stop received
FC00:172:16:254::2 03:12:35/03:12:35
```

To return to the example configuration, click [Example: Configuring BGP EVPN VXLAN with Dual Stack Underlay, on page 777](#).



CHAPTER 16

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 821](#)
- [Information About EVPN Microsegmentation, on page 821](#)
- [How to Configure EVPN Microsegmentation, on page 824](#)
- [Deployment of EVPN Microsegmentation, on page 824](#)
- [Configuration Example of EVPN Microsegmentation, on page 828](#)

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 73: 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 its network access by assigning an SGT tag.
- For configuration procedures to authenticate an endpoint at ingress, refer to the “Configuring Endpoint Admission Control” chapter in the *Cisco TrustSec Configuration Guide*.
 - 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 Policies” 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](#).

- **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 [Configuring EVPN VXLAN Integrated Routing and Bridging using Distributed Anycast 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 [Configuring EVPN VXLAN Integrated Routing and Bridging using Centralized Default Gateway](#).

- **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 gil/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 HundredGigE2/0/29
  switchport trunk allowed vlan 101,102,201,202
  switchport mode trunk
  cts manual
  policy static sgt 2 trusted
<snip: only relevant configuration shown>

```



CHAPTER 17

Configuring Additional BGP EVPN Features

This module discusses additional features of BGP EVPN such as Dynamic BGP Peering, and Route Map support.

- [Dynamic BGP Peering for EVPN, on page 829](#)
- [EVPN Route Map Support, on page 830](#)
- [BGP EVPN VRF Auto Route Distinguisher, on page 833](#)
- [BGP EVPN VRF Auto Route Target, on page 834](#)

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](#).

In Cisco IOS XE Release 17.11.1, support for Dynamic BGP Peering is extended to the L2VPN EVPN address family. To understand how to configure, verify, and troubleshoot BGP Dynamic Peering for the EVPN address family, refer to [Configure BGP DN for Multiple AF on Catalyst 9000 Series Switches](#).

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# <code>configure terminal</code>	Enters global configuration mode.
Step 2	route-map <i>route-map-name</i> Example:	Creates a route map and enters route-map configuration mode.

	Command or Action	Purpose
	Device(config)# route-map rmap1	
Step 3	match ip address prefix-list <i>prefix-list-name</i> Example: Device(config-route-map)# match ip address prefix-list plist1	Matches the BGP EVPN routes with the specified prefix-list.
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 <i>route-map-name</i> 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 <i>autonomous-system</i> 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 { <i>ip-address</i> <i>peer-group-name</i> } remote-as <i>number</i> 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 <i>ip-address</i> activate Example: Device(config-router-af)# neighbor 10.1.1.2 activate	Enables the exchange of information with a BGP neighbor.
Step 6	neighbor <i>ip-address</i> route-map <i>route-map-name</i> { in out } Example:	Applies a route map to incoming or outgoing routes.

	Command or Action	Purpose
	Device(config-router-af) # neighbor 10.1.1.2 route-map rmap1 in	
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 Cisco IOS XE Dublin 17.12.1, 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
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 VXLAN Network Identifier (VNI) 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]
```

The following table describes the different fields of the command syntax.

<i>vnid</i>	Specifies the VXLAN Network Identifier (VNI). The VNI range is from 1 to 2147483647.
evpn-instance	(Optional) Specify the evpn-instance keyword to autogenerate the stitching of Route Targets.

Example

```
VTEP(config)# router bgp 150
VTEP(config-bgp)# bgp router-id interface Loopback0
VTEP(config-bgp)# exit

VTEP(config)# vrf definition green
VTEP(config-vrf)# vnid 3000101 evpn-instance
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
```

In Cisco IOS XE Dublin 17.12.1, the **show vrf detail** command is enhanced to indicate if an RT has been assigned automatically.

Example

```
VTEP# show vrf detail green
VRF green (VRF Id = 182); default RD 10.1.10.1:121(auto); default VPNID <not set>
  New CLI format, supports multiple address-families
  vnid: 3000101 evpn-instance
  Flags: 0x380C
  Interfaces:
    Lo10                V1201                V1202
    V1203                V1204                V1205
```

```

Vl401                Vl402                Vl403
Vl404                Vl405                Vl2101
Address family ipv4 unicast (Table ID = 0xB6):
  Flags: 0x0
  Export VPN route-target communities
    RT:150:3000101 (auto)
  Import VPN route-target communities
    RT:150:3000101 (auto)
  Export VPN route-target stitching communities
    RT:150:3000101 (auto)
  Import VPN route-target stitching communities
    RT:150:3000101 (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 (Table ID = 0x1E000002):
  Flags: 0x0
  Export VPN route-target communities
    RT:150:3000101 (auto)
  Import VPN route-target communities
    RT:150:3000101 (auto)
  Export VPN route-target stitching communities
    RT:150:3000101 (auto)
  Import VPN route-target stitching communities
    RT:150:3000101 (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 ipv4 multicast not active
Address family ipv6 multicast not active

```



CHAPTER 18

BGP EVPN VXLAN Scalability Guide

- [Scale and Performance Capabilities for BGP EVPN VXLAN on Cisco Catalyst 9300 Series Switches, on page 837](#)

Scale and Performance Capabilities for BGP EVPN VXLAN on Cisco Catalyst 9300 Series Switches

This section provides information about the platform capabilities and the scale values for various components of a BGP EVPN VXLAN fabric. The numbers in the following table depend on the ternary content addressable memory (TCAM) scale. Use the SDM configuration templates to adjust the Layer 2 and Layer 3 TCAM forwarding table sizes based on the platform positioning. For more information, see *Configuring SDM Templates* module of the *System Management Configuration Guide* for the applicable release.

The scale values listed here are validated with a unidimensional configuration. The values provided in these tables focus on the scalability of one particular feature at a time.

The scale values are applicable to all the models, unless otherwise stated.

Product Specification							
Switch Model	C9300-L Series		C9300 Series	C9300 - UB/UBX Series		C9300-X Series	
BGP EVPN / VXLAN -- Leaf Scale							
SDM Template	Access	NAT	Access	Access	NAT	Access	NAT
VXLAN Leaf Node (Per Fabric Domain)	500	500	500	500	500	500	500
VRF	256	256	256	256	256	256	256
Layer 2 Virtual Network Instances (VNIs)	1024	1024	1024	1024	1024	750	750
Layer 3 VNIs	256	256	256	256	256	256	256
Layer 3 VRF SVI Interface	4094	4094	4094	4094	4094	4094	4094

Product Specification							
Switch Model	C9300-L Series		C9300 Series	C9300 - UB/UBX Series		C9300-X Series	
MAC Entries (Local L2 Network)	32000	32000	32000	64000	64000	32000	32000
MAC Entries (Remote EVPN VXLAN Network)	32000	32000	32000	32000	64000	32000	32000
Overlay IPv4 routes (LPM or Indirect Routes)	8000	8000	8000	8000	64000	64000	15000
Overlay IPv4 Host routes	32000	32000	32000	32000	48000	48000	39000
Overlay IPv6 routes (LPM or Indirect Routes)	4000	4000	4000	4000	32000	32000	7500
Overlay IPv6 Host routes	16000	16000	16000	16000	24000	24000	19500
Tenant Routed Multicast (TRM) IPv4	8000	8000	8000	8000	16000	16000	8000
TRM IPv6	4000	4000	4000	4000	8000	8000	4000
TRMv4: Overlay Multicast Routes (*,G and S,G)	8000	8000	8000	8000	16000	16000	8000
TRMv6: Overlay Multicast Routes (*,G and S,G)	4000	4000	4000	4000	8000	8000	4000
EVPN VXLAN Aware Flexible NetFlow - IPv4 Ingress Cache Entries	32000	32000	64000	64000	32000	32000	32000
EVPN VXLAN Overlay IPv4 Egress NetFlow Cache Entries (Per ASIC)	32000	32000	64000	64000	32000	32000	32000
EVPN VXLAN Overlay IPv6 Ingress NF Cache Entries (Per ASIC)	32000	32000	64000	64000	32000	32000	32000

Product Specification							
Switch Model	C9300-L Series		C9300 Series	C9300 - UB/UBX Series		C9300-X Series	
EVPN VXLAN Overlay IPv6 Egress NF Cache Entries (Per ASIC)	32000	32000	64000	64000	32000	32000	32000
Layer 2 VNI (L2VNI) Multicast Replication BUM Rate-Limiter	512	512	512	512	512	512	512
MicroSegmentation - Community VLAN to L2VNI	2000	2000	2000	2000	2000	2000	2000
NanoSegmentation - Isolated VLAN to L2VNI	384	384	384	384	384	384	384
Wide Area Bonjour (mDNS) over VXLAN Service Instance Count	15000	15000	15000	15000	15000	15000	15000
BGP EVPN / VXLAN -- Spine Scale							
SDM Template	Access	NAT	Access	Access	NAT	Access	NAT
BGP IPv4 Peer Scale	250	250	250	250	250	250	250
BGP IPv6 Peer Scale	250	250	250	250	250	250	250
BGP L2VPN EVPN Peer Scale	250	250	250	250	250	250	250
Overlay IPv4 routes (LPM or Indirect Routes)	8000	8000	8000	8000	64000	64000	15000
Overlay IPv4 Host routes	32000	32000	32000	32000	48000	48000	39000
Overlay IPv6 routes (LPM or Indirect Routes)	4000	4000	4000	4000	32000	32000	7500
Overlay IPv6 Host routes	16000	16000	16000	16000	24000	24000	19500
BGP EVPN / VXLAN -- Border Scale							
SDM Template	Access	NAT	Access	Access	NAT	Access	NAT

Product Specification							
Switch Model	C9300-L Series		C9300 Series	C9300 - UB/UBX Series		C9300-X Series	
EVPN to Layer 2 Handoff: IEEE 802.1Q	512	512	512	512	512	512	512
EVPN to Layer 2 Handoff: IEEE 802.1ad (QinQ)	512	512	512	512	512	512	512
EVPN to VRF Handoff: IP VRF (IPv4 and IPv6)	256	256	256	256	256	256	256
EVPN to MPLS Layer 3 VRF Unicast Handoff: VPNv4	256	256	256	256	256	256	256
EVPN to MPLS Layer 3 VRF Unicast Handoff: VPNv6	256	256	256	256	256	256	256
EVPN to MPLS Layer 3 VRF Multicast Handoff: mVPNv4	256	256	256	256	256	256	256
EVPN to MPLS Layer 3 VRF Multicast Handoff: mVPNv6	256	256	256	256	256	256	256
EVPN to VPLS Layer 2 Handoff: Virtual Forwarding Instances (VFIs)	512	512	512	512	512	512	512
EVPN to VPLS Layer 2 Handoff: Neighbors Per VFI	128	128	128	128	128	128	128
EVPN to VPLS Layer 2 Handoff: Pseudowire	512	512	512	512	512	512	512



CHAPTER 19

Troubleshooting BGP EVPN VXLAN

- [Troubleshooting Scenarios for BGP EVPN VXLAN, on page 841](#)
- [Troubleshooting Broadcast, Unknown Unicast, Multicast Traffic Forwarding, on page 842](#)
- [Troubleshooting Unicast Forwarding Between VTEPs in the Same VLAN Through a Layer 2 VNI, on page 846](#)
- [Troubleshooting Unicast Forwarding Between VTEPs in Different VLANs Through a Layer 3 VNI, on page 858](#)
- [Troubleshooting Unicast Forwarding Between a VXLAN Network and an IP Network, on page 871](#)
- [Troubleshooting Tenant Routed Multicast, on page 874](#)

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 74: EVPN VXLAN Topology



The following are the various troubleshooting scenarios that apply to BGP EVPN VXLAN for the topology illustrated in the [Figure 74: EVPN VXLAN Topology](#) 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, Unknown Unicast, Multicast Traffic Forwarding

This scenario might occur when host device 2 attempts to learn the ARP for host device 3 in [Figure 74: EVPN VXLAN Topology](#), on page 842. Perform the checks listed in the following table before troubleshooting BUM traffic forwarding:

Table 66: Scenario 1: Broadcast, Unknown 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.
Are the hosts in the same subnet or in different subnets?	Perform any of the following steps: <ul style="list-style-type: none"> • 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 macTable vlan <i>vlan-id</i> command in privileged EXEC mode on the local VTEP and check if the MAC address of the remote host device is displayed 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 843](#)
2. [Set Up Embedded Capture Towards the Core-Facing Interface, on page 844](#)
3. [Ping the Remote Host Device, on page 844](#)
4. [Verify that an ARP Request Has Been Received and a Multicast Route Has Been Built, on page 844](#)
5. [Confirm the Presence of ARP Request Replies in Embedded Capture, on page 845](#)
6. [Verify that the Encapsulated ARP Request is Leaving in a Multicast Group to a VXLAN UDP Destination Port, on page 845](#)
7. [Verify that the ARP Reply from Core Interface is Encapsulated in Unicast to a VXLAN UDP Destination Port, on page 846](#)

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 Tw1/0/1  <<- 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 Tw1/0/1:10
<<- 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 TwoGigabitEthernet1/0/2 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:
TwoGigabitEthernet1/0/2, 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: TwoGigabitEthernet1/0/2, 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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000	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
2	0.000	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

```

> Frame 1: 110 bytes on wire (880 bits), 110 bytes captured (880 bits) on interface 0
v Ethernet II, Src: 74:a2:e6:4f:c9:00, Dst: 01:00:5e:0a:0a:0a
  > Destination: 01:00:5e:0a:0a:0a
  > Source: 74:a2:e6:4f:c9:00
  Type: IPv4 (0x0800)
  > Internet Protocol Version 4, Src: 10.255.1.1, Dst: 239.10.10.10
  v User Datagram Protocol, Src Port: 65419 (65419), Dst Port: 4789 (4789)
    Source Port: 65419
    Destination Port: 4789
    Length: 76
    > Checksum: 0x0000 (none)
    [Stream index: 0]
  v Virtual eXtensible Local Area Network
    > Flags: 0x0000, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 10001
    Reserved: 0
  v Ethernet II, Src: 00:5f:86:02:10:c6, Dst: ff:ff:ff:ff:ff:ff
    > Destination: ff:ff:ff:ff:ff:ff
    > Source: 00:5f:86:02:10:c6
    Type: ARP (0x0806)
    Trailer: 0000000000000000000000000000000000000000000000000000000000000000
    Address Resolution Protocol (request)
  
```

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.

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000	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
2	0.000	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

```

> Frame 2: 110 bytes on wire (880 bits), 110 bytes captured (880 bits) on interface 0
v Ethernet II, Src: 74:a2:e6:4f:c9:00, Dst: 70:35:09:56:7e:d6
  > Destination: 70:35:09:56:7e:d6
  > Source: 74:a2:e6:4f:c9:00
  Type: IPv4 (0x0800)
  > Internet Protocol Version 4, Src: 10.255.1.2, Dst: 10.255.1.1
  v User Datagram Protocol, Src Port: 65350 (65350), Dst Port: 4789 (4789)
    Source Port: 65350
    Destination Port: 4789
    Length: 76
    > Checksum: 0x0000 (none)
    [Stream index: 1]
  v Virtual eXtensible Local Area Network
    > Flags: 0x0000, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 10001
    Reserved: 0
  v Ethernet II, Src: 28:52:61:bf:a9:46, Dst: 00:5f:86:02:10:c6
    > Destination: 00:5f:86:02:10:c6
    > Source: 28:52:61:bf:a9:46
    Type: ARP (0x0806)
    Trailer: 0000000000000000000000000000000000000000000000000000000000000000
    Address Resolution Protocol (reply)
  
```

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 67: 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 847](#)
2. [Ensure that an NVE Peer is Present for the Layer 2 VNI, on page 849](#)
3. [Verify the Provisioning of the Layer 2 VNI in NVE Component, on page 849](#)
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 850](#)

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:
  TwoGigabitEthernet1/0/1 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   Tw1/0/1:10      <<- 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 nvl: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 nvl: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 75: 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 852
2. Verify that Both Local and Remote MAC Addresses are Learned in FED-MATM, on page 852

3. Confirm that the ICMP Echo Request Leaves VTEP 1 Encapsulated and Goes to a UDP Destination Port on VTEP 2, on page 853
4. Verify ARP for Local Host Devices, on page 853
5. Verify that the MAC Address Entries are Learned in SISF Device Tracking Table, on page 853
6. Verify that EVPN Manager Has Been Updated with the MAC Address Entries, on page 854
7. Verify that EVPN Manager Has Updated the MAC Routes into Layer 2 RIB, on page 855
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 855
9. Verify that the MAC Routes Learned from BGP and Updated to Layer 2 RIB are Also Updated to L2FIB, on page 857



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 tw 1/0/1 vlan 10
      Mac Address Table
-----
Vlan    Mac Address      Type    Ports
----    -
10      005f.8602.10c6   DYNAMIC Tw1/0/1    <<- 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    TwoGigabitEthernet1/0/1

<<- 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      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    GigabitEthernet2/0/1

<<- 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:

→	1	0.000	10.10.10.11	10.10.10.12	ICMP	164	Echo (ping) request
←	2	0.000	10.10.10.12	10.10.10.11	ICMP	164	Echo (ping) reply

```

▶ Frame 1: 164 bytes on wire (1312 bits), 164 bytes captured (1312 bits) on interface 0
▶ Ethernet II, Src: 00:00:00:00:00:00, Dst: 00:00:00:00:00:00
▶ Internet Protocol Version 4, Src: 10.255.1.1, Dst: 10.255.1.2 ← Lo999 VTEP loopbacks
▶ User Datagram Protocol, Src Port: 65419 (65419), Dst Port: 4789 (4789)
▼ Virtual eXtensible Local Area Network
  ▶ Flags: 0x0800, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 10001 ← L2 VNI 10001 Vlan 10
    Reserved: 0
  ▶ Ethernet II, Src: 00:5f:86:02:10:c6, Dst: 28:52:61:bf:a9:46 ← Native Source/Dest IP/MAC
  ▶ Internet Protocol Version 4, Src: 10.10.10.11, Dst: 10.10.10.12 ← Native Source/Dest IP/MAC
  ▶ Internet Control Message Protocol

```

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    Tw1/0/1        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          Tw1/0/1:10
<<- 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
Ethernet Tag ID:     0
Next Hop(s):         TwoGigabitEthernet1/0/1 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
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 address
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      Tw1/0/1:10   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 to 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 12fib 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   Tw1/0/1:10
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 68: 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.
Do you have an SVI interface configured for the remote subnet?	Run the show ip interface brief exclude unassigned command in privileged EXEC mode on the VTEP.

Check to be Performed	Steps to Follow
Do you have the EVPN instance configured on your local VTEP?	Run the following commands in privileged EXEC mode on the VTEP: <ul style="list-style-type: none"> • 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 859](#)
2. [Verify that the IP VRF is Provisioned with the Correct SVIs, Stitching Route-Targets, and Route Distinguisher, on page 860](#)
3. [Verify that Both Layer 2 and Layer 3 VNIs are provisioned in the VRF and are UP, on page 861](#)
4. [Verify that EVPN Manager is Updated from the NVE with all the Layer 2 and IRB Attributes, on page 862](#)
5. [Verify that the Remote Layer 3 VNI Details are Learned on Each VTEP, on page 863](#)
6. [Verify that the Layer 3 VNI Tunnel Pseudoport is Installed into Layer 2 FIB in the Core VLAN, on page 863](#)

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       V110
                                      V112
                                      V199
```

```
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:
V110 V112 V199
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       V110
                                      V113
                                      V199
```

```
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  Tw1/0/1:12
<<- 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:
  TwoGigabitEthernet1/0/1 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 76: 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 865
2. Verify that MAC Address and IP Address Entries are Learned in EVPN Manager, on page 866
3. Verify that MAC Address and IP Address Entries are Learned in Layer 2 RIB, on page 867
4. Verify that Local MAC Address and IP Address Entries are Learned in MAC VRF, on page 867
5. Verify that Remote MAC-IP Address Pair is Learend in the VRF, on page 868
6. Verify that IP Routes are Inserted in RIB, on page 869
7. Verify that the Adjacency Table Contains Entries for the VRF-Enabled Core VLAN Interface, on page 869
8. Confirm that Adjacency Exists to the VTEP Tunnel IP Address for a Host Device in IP VRF, on page 870
9. Confirm that Adjacency Exists to Reach Tunnel Destination, on page 870
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 870

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 state	Time left	Link Layer Address	Interface	vlan	prlvl	age
ARP 10.12.12.12		005f.8602.10e7	Tw1/0/1	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 state	Time left	Link Layer Address	Interface	vlan	prlvl	age
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	Tw1/0/1:12

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:           TwoGigabitEthernet1/0/1 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      Tw1/0/1:12
```

```
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      Gi2/0/1:13
```

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 Learned 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 TwoGigabitEthernet1/0/2
    
```

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:

Seq	Len	Source	Destination	Protocol	Length	Operation
→ 3	0.000	10.12.12.12	10.13.13.13	ICMP	164	Echo (ping) request
← 4	0.000	10.13.13.13	10.12.12.12	ICMP	164	Echo (ping) reply
→ 5	0.000	10.12.12.12	10.13.13.13	ICMP	164	Echo (ping) request
← 6	0.000	10.13.13.13	10.12.12.12	ICMP	164	Echo (ping) reply

```

▶ Frame 3: 164 bytes on wire (1312 bits), 164 bytes captured (1312 bits) on interface 0
▶ Ethernet II, Src: 00:00:00:00:00:00, Dst: 00:00:00:00:00:00
▶ Internet Protocol Version 4, Src: 10.255.1.1, Dst: 10.255.2.1 ← Tunnel Endpoint IPs
▶ User Datagram Protocol, Src Port: 65478 (65478), Dst Port: 4789 (4789)
▼ Virtual eXtensible Local Area Network
  ▶ Flags: 0x0800, VXLAN Network ID (VNI)
    Group Policy ID: 0
    VXLAN Network Identifier (VNI): 9999 ← L3 VNI 9999 VRF vxlan
    Reserved: 0
  ▶ Ethernet II, Src: 00:01:00:01:00:00, Dst: 74:86:0b:c4:b7:5d ← VTEP-2 Dst: RMAC
  ▶ Internet Protocol Version 4, Src: 10.12.12.12, Dst: 10.13.13.13
  ▶ Internet Control Message Protocol
    
```

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 69: 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?	<p>Check the local subnets (or the SVI interfaces) if the remote subnet is present.</p> <p>Note Local subnet has the remote subnet listed even in the case of scenario 3.</p>
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

See [Verifying the Provisioning of an EVPN VXLAN Layer 3 Overlay Network, on page 859](#) for detailed steps.

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 871](#)
2. [Check the Table Entries for BGP, EVPN, and VPNv4 Tables, on page 871](#)
3. [Confirm that Adjacency exists to Reach Tunnel Destination, on page 874](#)

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  ports  siHandle
      riHandle                diHandle                *a_time *e_time
-----
99    7035.0956.7edd             0x8002  0      0      64  0x7ffa48d61be8  0x7ffa48d630b8
      0x0                      0x5154                0
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 TwoGigabitEthernet1/0/2
```

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

Integration with Cisco DNA Service for Bonjour

- [Cisco DNA Service for Bonjour Solution Overview, on page 877](#)
- [Configuring Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 889](#)
- [Configuring VRF-Aware Local Area Bonjour Services, on page 911](#)



CHAPTER 20

Cisco DNA Service for Bonjour Solution Overview

- [About the Cisco DNA Service for Bonjour Solution, on page 877](#)
- [Solution Components, on page 878](#)
- [Supported Platforms, on page 879](#)
- [Supported Network Design, on page 880](#)


About the Cisco DNA Service for Bonjour Solution

The Apple Bonjour protocol is a zero-configuration solution that simplifies rich services and enables intuitive experience between connected devices, services, and applications. Using Bonjour, you can discover and use IT-managed, peer-to-peer, audio and video, or Internet of Things (IoT) services with minimal intervention and technical knowledge. Bonjour is originally designed for single Layer 2 small to mid-size networks, such as home or branch networks. The Cisco DNA Service for Bonjour solution eliminates the single Layer 2 domain constraint and expands the matrix to enterprise-grade traditional wired and wireless networks, including overlay networks such as Cisco Software-Defined Access (SD-Access) and industry-standard BGP EVPN with VXLAN. The Cisco Catalyst 9000 Series LAN switches, Cisco Nexus 9300 Series Switches, and Cisco Catalyst 9800 Series Wireless Controller follow the industry standard, RFC 6762-based multicast DNS (mDNS) specification to support interoperability with various compatible wired and wireless consumer products in enterprise networks.

The Cisco Wide Area Bonjour application on Cisco DNA Center enables mDNS service routing to advertise and discover services across enterprise-grade wired and wireless networks. The new-distributed architecture is designed to eliminate mDNS flood boundaries and transition to unicast-based service routing, providing policy enforcement points and enabling the management of Bonjour services.

The following figure illustrates how the Cisco Wide Area Bonjour application operates across two integrated service-routing domains.

Figure 77: Cisco Wide Area Bonjour Solution Architecture



- Local Area Service Discovery Gateway Domain - Unicast Mode:** The new enhanced Layer 2 unicast policy-based deployment model. The new mDNS service discovery and distribution using the Layer 2 unicast address enables flood-free LAN and wireless networks. Cisco Catalyst 9000 Series Switches and Cisco Catalyst 9800 Series Wireless Controller in Layer 2 mode introduce a new service-peer role, replacing the classic flood-n-learn, for new unicast-based service routing support in the network. The service-peer switch and wireless controller also replace mDNS flood-n-learn with unicast-based communication with any RFC 6762 mDNS-compatible wired and wireless endpoints.
- Wide-Area Service Discovery Gateway Domain:** The Wide Area Bonjour domain is a controller-based solution. The Bonjour gateway role and responsibilities of Cisco Catalyst and Cisco Nexus 9300 Series Switches are extended from a single SDG switch to an SDG agent, enabling Wide Area Bonjour service routing beyond a single IP gateway. The network-wide distributed SDG agent devices establish a lightweight, stateful, and reliable communication channel with a centralized Cisco DNA Center controller running the Cisco Wide Area Bonjour application. The SDG agents route locally discovered services based on the export policy.



Note The classic Layer 2 multicast flood-n-learn continues to be supported on wired and wireless networks with certain restrictions to support enhanced security and location-based policy enforcement. The Cisco Catalyst and Cisco Nexus 9300 Series Switches at Layer 3 boundary function as an SDG to discover and distribute services between local wired or wireless VLANs based on applied policies.

Solution Components

The Cisco DNA Service for Bonjour solution is an end-to-end solution that includes the following key components and system roles to enable unicast-based service routing across the local area and Wide Area Bonjour domain:

- Cisco Service Peer:** Cisco Catalyst Switches and Cisco Wireless Controllers in Layer 2 access function in service peer mode to support unicast-based communication with local attached endpoints and export service information to the upstream Cisco Catalyst SDG agent in the distribution layer.



Note Cisco Nexus 9300 Series Switches don't support unicast-based service routing with downstream Layer 2 access network devices.

- **Cisco SDG Agent:** Cisco Catalyst and Cisco Nexus 9300 Series Switches function as an SDG agent and communicate with the Bonjour service endpoints in Layer 3 access mode. At the distribution layer, the SDG agent aggregates information from the downstream Cisco service peer switch and wireless controller, or local Layer 2 networks, and exports information to the central Cisco DNA controller.



Note Cisco Nexus 9300 Series Switches don't support multilayer LAN-unicast deployment mode.

- **Cisco DNA controller:** The Cisco DNA controller builds the Wide Area Bonjour domain with network-wide and distributed trusted SDG agents using a secure communication channel for centralized services management and controlled service routing.
- **Endpoints:** A Bonjour endpoint is any device that advertises or queries Bonjour services conforming to RFC 6762. The Bonjour endpoints can be in either LANs or WLANs. The Cisco Wide Area Bonjour application is designed to integrate with RFC 6762-compliant Bonjour services, including AirPlay, Google Chrome cast, AirPrint, and so on.

Supported Platforms

The following table lists the supported controllers, along with the supported hardware and software versions.

Table 70: Supported Controllers with Supported Hardware and Software Versions

Supported Controller	Hardware	Software Version
Cisco DNA Center appliance	DN2-HW-APL DN2-HW-APL-L DN2-HW-APL-XL	Cisco DNA Center, Release 2.3.6
Cisco Wide Area Bonjour application	—	2.4.660.75403

The following table lists the supported SDG agents along with their licenses and software requirements.

Table 71: Supported SDG Agents with Supported License and Software Requirements

Supported Platform	Supported Role	Local Area SDG	Wide Area SDG	Minimum Software
Cisco Catalyst 9200 Series Switches	SDG agent Service peer	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1

Supported Platform	Supported Role	Local Area SDG	Wide Area SDG	Minimum Software
Cisco Catalyst 9200L Series Switches	SDG agent Service peer	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9300 and 9300-X Series Switches	Service peer SDG agent	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9400 and 9400-X Series Switches	Service peer SDG agent	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9500 and 9500-X Series Switches	Service peer SDG agent	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9500 High Performance Series Switches	Service peer SDG agent	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9600 and 9600-X Series Switches	Service peer SDG agent	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9800 Wireless Controller	Service peer	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Catalyst 9800-L Wireless Controller	Service peer	Cisco DNA Advantage	Cisco DNA Advantage	Cisco IOS XE Release 17.11.1
Cisco Nexus 9300 Series Switches	SDG agent	Cisco DNA Advantage	Cisco DNA Advantage	Cisco NX-OS Release 10.2(3)F

Supported Network Design


The Cisco DNA Service for Bonjour supports a broad range of enterprise-grade networks. The end-to-end unicast-based Bonjour service routing is supported on traditional, Cisco SD-Access, and BGP EVPN-enabled wired and wireless networks.

Traditional Wired and Wireless Networks

Traditional networks are classic Layer 2 or Layer 3 networks for wired and wireless modes deployed in enterprise networks. Cisco DNA Service for Bonjour supports a broad range of network designs to enable end-to-end service routing and replace flood-n-learn-based deployment with a unicast mode-based solution.

The following figure illustrates traditional LAN and central-switching wireless local mode network designs that are commonly deployed in an enterprise.

Figure 78: Enterprise Traditional LAN and Wireless Local Mode Network Design




357153

Wired Networks

The following figure shows the supported traditional LAN network designs that are commonly deployed in an enterprise.

Figure 79: Enterprise Wired Multilayer and Routed Access Network Design



The Cisco Catalyst or Cisco Nexus 9300 Series Switches in SDG agent role that provide Bonjour gateway functions are typically IP gateways for wired endpoints that could reside in the distribution layer in multilayer network designs, or in the access layer in Layer 3 routed access network designs:


- **Multilayer LAN—Unicast Mode:** In this deployment mode, the Layer 2 access switch provides the first-hop mDNS gateway function to locally attached wired endpoints. In unicast mode, the mDNS services are routed to the distribution layer systems providing IP gateway and SDG agent mode. The policy-based service routing between the SDG agents is performed by the Cisco DNA Center controller.
- **Multilayer LAN—Flood-n-Learn Mode:** In this deployment mode, the Layer 2 access switch or wireless controller are in mDNS passthrough modes with the Cisco Catalyst or Cisco Nexus 9300 Series Switches operating in the SDG agent mode. The mDNS gateway function at distribution layer in a network enables inter-VLAN mDNS local proxy. It also builds stateful Wide Area Bonjour unicast service routing with the Cisco DNA Center to discover or distribute mDNS services beyond a single IP gateway.
- **Routed Access:** In this deployment mode, the first-hop Cisco Catalyst or Cisco Nexus 9300 Series Switch is an IP gateway boundary and, therefore, it must also perform the SDG agent role. The policy-based service routing between the SDG agents is performed by the Cisco DNA Center controller.

Wireless Networks

The Cisco DNA Service for Bonjour extends the single wireless controller mDNS gateway function into the Wide Area Bonjour solution. The mDNS gateway on Cisco Catalyst 9800 Series Wireless Controller can be deployed in an enhanced mode as a service peer. In this mode, the wireless controller builds unicast service routing with an upstream Cisco Catalyst gateway switch for end-to-end mDNS service discovery. It replaces the classic flood-n-learn mDNS services from wired network using mDNS AP or other methods.

The following figure shows the supported traditional wireless LAN network designs that are commonly deployed in an enterprise. Based on the wireless network design, the mDNS gateway function may be on the wireless controller, or first-hop Layer 2 or Layer 3 Ethernet switch of an Access Point in local-switching mode.

Figure 80: Enterprise Traditional Wireless LAN Network Design



The Cisco DNA Service for Bonjour supports the following modes for wireless LAN networks:

- Local Mode:** In the central switching wireless deployment mode, the m-DNS traffic from local mode Cisco access points is terminated on the Cisco Catalyst 9800 Series Wireless Controller. The Cisco Catalyst 9800 Series Wireless Controller extends the mDNS gateway function to the new service peer mode. The wireless controller can discover and distribute services to local wireless users and perform unicast service routing over a wireless management interface to the upstream Cisco Catalyst Switch in the distribution layer, which acts as the IP gateway and the SDG agent.
- FlexConnect—Central:** The mDNS gateway function for Cisco access point in FlexConnect central switch SSID functions consistently as described in **Local Mode**. The new extended mDNS gateway mode on the Cisco Wireless Controller and upstream service routing with SDG agent operate consistently to discover services across network based on policies and locations.
- FlexConnect—Local:** In FlexConnect local switching mode, the Layer 2 access switch in mDNS gateway service peer mode provides the policy-based mDNS gateway function to locally attached wired and wireless users. The Cisco Catalyst Switches in the distribution layer function as SDG agents and enable


mDNS service-routing across all Layer 2 ethernet switches to support unicast-based service routing to LAN and wireless LAN user groups.

- **Embedded Wireless Controller—Access Point:** The Layer 2 access switch in service peer mode provides unified mDNS gateway function to wired and wireless endpoints associated with Cisco Embedded Wireless Controller on Cisco Catalyst 9100 Series Access Points. The SDG agent in the distribution layer provides unicast service routing across all Layer 2 service peer switches in the Layer 2 network block without any mDNS flooding.

Cisco SD-Access Wired and Wireless Networks


Cisco SD-Access-enabled wired and wireless networks support Cisco DNA Service for Bonjour across fabric networks. The Cisco Catalyst 9000 Series Switches support VRF-aware Wide Area Bonjour service routing to provide secure and segmented mDNS service discovery and distribution management for virtual networks. The VRF-aware unicast service routing eliminates the need to extend Layer 2 flooding, and improves the scale and performance of the fabric core network and endpoints.

Figure 81: Cisco SD-Access Wired and Wireless Network Design



Cisco SD-Access supports flexible wired and wireless network design alternatives to manage fully distributed, integrated, and backward-compatible traditional network infrastructure. Wide Area Bonjour service routing is supported in all network designs providing intuitive user experience. The following figure illustrates the various SD-Access enabled wired and wireless network design alternatives.

Figure 82: Cisco SD-Access Wired and Wireless Network Design Alternatives



The Cisco DNA Service for Bonjour for SD-Access enabled wired and fabric, or traditional mode-wireless networks use two-tier service routing providing end-to-end unicast-based mDNS solution. Based on the network design, each solution component is enabled in a unique role to support the Wide Area Bonjour domain:

- Fabric Edge SDG Agent:** The Layer 3 Cisco Catalyst Fabric Edge switch in the access layer configured as SDG agent provides unicast-based mDNS gateway function to the locally attached wired and wireless endpoints. The VRF-aware mDNS service policy provides network service security and segmentation in a virtual network environment. The mDNS services can be locally distributed and routed through centralized Cisco DNA Center.
- Policy Extended Node:** The Layer 2 Cisco Catalyst access layer switch enables first-hop mDNS gateway function without flooding across the Layer 2 broadcast domain. The unicast-based service routing with upstream Fabric Edge switch in the distribution layer enables mDNS service routing within the same Layer 2 network block. It can also perform remote service discovery and distribution from centralized Cisco DNA Center.
- Cisco Wireless Controller:** Based on the following wireless deployment modes, Cisco Wireless Controller supports unique function to enable mDNS service routing in Cisco SD-Access enabled network:
 - Fabric-Enabled Wireless:** Cisco Wireless Controller doesn't require any mDNS gateway capability to be enabled in distributed fabric-enabled wireless deployments.
 - Local Mode Wireless:** As Cisco Wireless Controller provides central control and data plane termination, it provides mDNS gateway in service peer mode for wireless endpoints. The wireless controller provides mDNS gateway between locally associated wireless clients. The wireless controller builds service routing with upstream SDG agent Catalyst switch providing IP gateway and service routing function for wireless endpoints.
 - Embedded Wireless Controller—Switch:** The Cisco Embedded Wireless Controller solution enables the lightweight integrated wireless controller function within the Cisco Catalyst 9300 Series

Switch. The Cisco Catalyst switches in the distribution layer function as SDG agents to the wired and wireless endpoints. The SDG agent in the distribution layer provides unicast service routing across all wireless access points and Layer 2 service peer switches without mDNS flooding.

- **Cisco DNA Center Controller:** The Cisco Wide Area Bonjour application on Cisco DNA Center supports policy and location-based service discovery, and distribution between network-wide distributed Fabric Edge switches in SDG agent mode.


The Wide Area Bonjour communication between the SDG agent and controller takes place through the network underlay. Based on policies, the SDG agent forwards the endpoint announcements or queries to the Cisco DNA Center. After discovering a service, the endpoints can establish direct unicast communication through the fabric overlay in the same virtual network. The inter-virtual network unicast communication takes place through the Fusion router or external Firewall system. This communication is subject to the configured overlay IP routing and Security Group Tag (SGT) policies.

BGP EVPN Networks

The BGP EVPN-based technology provides a flexible Layer 3 segmentation and Layer 2 extension overlay network. The VRF and EVPN VXLAN-aware Wide Area Bonjour service routing provides secure and segmented mDNS service solution. The overlay networks eliminate mDNS flooding over EVPN-enabled Layer 2 extended networks and solve the service reachability challenges for Layer 3 segmented routed networks in the fabric.

The following figure shows the BGP EVPN leaf switch in the distribution layer, supporting overlay Bonjour service routing for a BGP EVPN-enabled traditional Layer 2 wired access switch and traditional wireless local mode enterprise network interconnected through various types of Layer 2 networks and Layer 3 segmented VRF-enabled networks.


Figure 83: Overlay Bonjour Service for a BGP EVPN-Enabled Enterprise Network



Cisco DNA Service for Bonjour supports all the industry-standard overlay network designs enabling end-to-end unicast-based mDNS service routing, and preventing flooding and service boundary limitation across wired and wireless networks.

The following figure illustrates the various BGP EVPN VXLAN reference overlay network design alternatives. This network design enables end-to-end mDNS service discovery and distribution based on overlay network policies.

Figure 84: BGP EVPN VXLAN Wired and Wireless Design Alternatives



The Cisco Catalyst and Cisco Nexus 9000 Series Switches can be deployed in Layer 2 or Layer 3 leaf roles supporting mDNS service routing for a broad range of overlay networks. In any role, the mDNS communication is limited locally and supports end-to-end unicast-based service routing across Wide Area Bonjour domain:

- **Layer 2 Leaf SDG Agent:** The Cisco Catalyst or Cisco Nexus switches can be deployed as Layer 2 leaf supporting end-to-end bridged network with IP gateway within or beyond BGP EVPN VXLAN fabric network. By default, the mDNS is flooded as Broadcast, Unknown Unicast, Multicast (BUM) over the fabric-enabled core network. This mDNS flooding may impact network performance and security. The Layer 2 leaf, enabled as SDG agent, prevents mDNS flooding over VXLAN and supports unicast-based service routing.
- **Layer 3 Leaf SDG Agent:** The Cisco Catalyst or Cisco Nexus switches can be deployed as SDG agent supporting Layer 3 overlay network in BGP EVPN VXLAN fabric. The IP gateway and mDNS service boundary is terminated at the SDG agent switches and remote services can be discovered or distributed through centralized Cisco DNA Center.
- **Local Mode Wireless:** The centralized wireless local mode network can be terminated within or outside the EVPN VXLAN fabric domain to retain network segmentation and service discovery for wireless endpoints. The Cisco Catalyst 9800 Series Wireless Controller in service peer mode can build unicast service routing with distribution layer IP and SDG agent Cisco Catalyst switch to discover services from BGP EVPN VXLAN fabric overlay network.
- **Cisco DNA Center:** Cisco DNA Center supports Wide Area Bonjour capability to dynamically discover and distribute mDNS services based on Layer 2 or Layer 3 Virtual Network ID (VNID) policies to route the mDNS services between SDG agent switches in the network.

For more information about BGP EVPN networks, see [Cisco DNA Service for Bonjour Configuration Guide, Cisco IOS XE Bengaluru 17.6.x \(Catalyst 9600 Switches\)](#).



CHAPTER 21

Configuring Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks

- [Prerequisites for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 889](#)
- [Restrictions for Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 890](#)
- [Information About Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 890](#)
- [How to Configure Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 893](#)
- [Verifying Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 902](#)
- [Additional References for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks, on page 909](#)

Prerequisites for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks

This section provides the list of prerequisites for a Cisco Catalyst leaf switch that needs to be deployed in SDG Agent mode.:

- Ensure that you successfully configure and operate the BGP EVPN VXLAN overlay networks on the Cisco Catalyst devices before you configure Cisco Local Area and Wide Area Bonjour for LAN and WLAN networks.
- Verify that the targeted leaf switch is supported in SDG Agent and the Layer 2 access switch is supported in Service-Peer mode. See [Supported Platforms, on page 879](#) for more information.
- Verify that the targeted SDG Agent leaf switch, Service-Peer switch and Service-Peer wireless controller (WLC) run on the minimum required Cisco IOS XE software version.
- Ensure that the SDG Agent leaf switch, Service-Peer switch, and Service-Peer WLC run on a valid Cisco DNA Advantage license.
- In a multilayer network with Layer 2 unicast service-routing between an SDG Agent leaf switch in distribution layer and Service-Peer, ensure that the connection is through a Layer 2 trunk in static mode.
- Ensure that the Cisco DNA Center has IP connectivity with the SDG Agent leaf switch in either the underlay or overlay network.

- Ensure that the SDG Agent leaf switch has IP connectivity with the Service-Peer switch and Service-Peer WLC in the same IPv4 subnet when the traffic is globally routed through the management VLAN.
- Verify that wireless AP multicast is configured in the underlay network and Cisco Wireless APs have successfully joined the AP multicast group announced by the WLC.



Note The leaf switch can also enable unicast-based service-routing with a downstream Layer 2 access switch and Catalyst 9800 Series WLC.

Restrictions for Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks

- Only Cisco Catalyst 9000 Series Layer 2 Ethernet switches that are connected to a distribution layer leaf switch can be deployed in Service-Peer mode. No other classic Cisco Catalyst LAN switches are supported in Service-Peer mode.
- Cisco Embedded WLC is not supported on Cisco Catalyst Series switches.
- Multicast DNS (mDNS) Trust port is not supported on the Layer 2 port channel interface of Cisco Catalyst Series switches.
- Local-area service-routing is not supported on the management port of Cisco Catalyst Series switches.
- In a multilayer topology wherein the endpoints are connected to the Service-Peer and the SDG Agent deploys Layer 3 Virtual Network Instance (L3VNI), the mDNS services are not cached at the SDG Agent. The services are only cached at the Service-Peer and advertised to the SDG Agent.


L3VNI deployment is applicable in a routed-access topology wherein the endpoints are connected to the SDG Agent and the services are advertised to the controller.

Information About Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks

Cisco Catalyst 9000 Series leaf switches or VTEPs introduce hierarchical service-routing for standard-based BGP EVPN VXLAN overlay networks. Leaf switches can function as Service Discovery Gateway Agents (SDG-Agents) to enable mDNS service-routing for Layer 3 and Layer 2 overlay networks. The enhanced gateway function (at the first hop of wired and wireless networks) communicates with directly associated, industry-standard, [RFC 6762](#)-compliant multicast DNS (mDNS) end points in Layer 2 unicast mode. Communication in unicast mode eliminates Layer 2 mDNS traffic flooding and enables service discovery across Layer 3 networks for large scale and enterprise-grade LAN and WLAN overlay networks. By doing this, unicast mode enhances security, bandwidth, scale, and performance for underlay and overlay fabric networks.

The following image illustrates a BGP EVPN VXLAN fabric that supports service-routing for wired and wireless user endpoints over Layer 3 overlay networks:

Figure 85: Cisco Wide Area Bonjour over BGP EVPN VXLAN Layer 3 Overlay Networks



A leaf switch in enterprise campus access or distribution layer can perform Wide Area Bonjour service-routing. Service-routing allows the leaf switch to establish stateful and reliable communication with a centralized Cisco DNA Center in the underlay network. The leaf can also discover and distribute policy-based services. Service data communication between mDNS sources and receivers over an EVPN VXLAN network uses segmented Layer 3 overlay forwarding topologies. A leaf switch can extend unicast-based service-routing through a downstream Layer 2 access device to build a Local Area Bonjour domain for the VXLAN fabric. The domain allows end-to-end, unicast-based service-routing for EVPN VXLAN Layer 3 overlay networks. The downstream Layer 2 access device in the fabric can be one of:

- A Cisco Catalyst 9000 Series Ethernet switch
- A Cisco Catalyst 9800 Series WLC

Understanding Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks

A BGP EVPN Layer 3 leaf switch provides an IP gateway to a Layer 3-segmented, VRF-enabled interface to attach wired and wireless endpoints locally. Each BGP EVPN leaf switch in an access or distribution layer network provides a demarcation point to the broadcast boundary of its local Layer 2 network. This demarcation prevents mDNS service discovery and distribution beyond the leaf switch in the access or distribution layer network.


Cisco Wide Area Bonjour for BGP EVPN overlay networks enables unicast-based mDNS service discovery and distribution beyond the boundary of the single IP network. A network-wide distributed BGP EVPN leaf switch in an access or distribution layer network builds stateful and reliable communication with a Cisco DNA Center. This link enables policy- and location-based mDNS service-routing between the leaf switches for Layer 3 overlay networks.

The BGP EVPN fabric boundary can be initiated from the access or distribution layer network depending on network design. The following [Figure 86: Wide Area Bonjour for Wired and Wireless BGP EVPN Layer 3 Overlay Networks](#) provides two deployment alternatives to determine the handling of mDNS communication with downstream Layer 2 switches or WLCs and upstream fabric networks:

- **Leaf Access and Distribution:** The LAN access in Layer 3 mode provides the Layer 2–Layer 3 boundary to enable mDNS service-routing over the EVPN fabric network. The leaf switch in distribution layer can connect to a downstream Layer 2 WLC in Service-Peer mode and enable Unicast-based service-routing.

- **Leaf Distribution:** The LAN and Wireless distribution layer provides the Layer 2–Layer 3 boundary to enable mDNS service-routing over the EVPN fabric network. The leaf switch in Distribution layer can connect to a downstream Layer 2 LAN access switch and WLC in Service-Peer mode and enable Unicast-based service-routing.

Figure 86: Wide Area Bonjour for Wired and Wireless BGP EVPN Layer 3 Overlay Networks



mDNS service-routing for EVPN VXLAN Layer 3 overlay networks supports dynamic service discovery from a local network and performs service-routing with a centralized Cisco DNA-Center. The mDNS service-routing functions in the following ways with Bonjour domains:

- **Local Area Bonjour Domain:** A BGP EVPN Layer 3 leaf switch discovers and distributes mDNS services dynamically. The leaf switch distributes the mDNS services to directly or indirectly attached wired and wireless endpoints in a Layer 2 network. The leaf switch in distribution layer can establish unicast-based service-routing with a downstream Layer 2 access device to improve scale, performance, and location-based service capabilities. The downstream Layer 2 access device can be either a Cisco Catalyst Series switch or a Cisco 9800 Series WLC.

See [Configuring Local Area Bonjour Unicast Mode for Wired and Wireless Local Mode Networks](#) for detailed configuration steps.

- **Wide Area Bonjour Domain:** A BGP EVPN Layer 3 leaf switch and Wide Area Bonjour application in a Cisco DNA Center work together to build secure service-routing and peering. The link-up enables policy-based and location-based mDNS service-routing across Layer 3 overlay networks. The Layer 3 overlay network allows the exchange of VXLAN data only between the leaf switch and the Cisco DNA Center. The endpoints of mDNS service providers and receivers do not forward VXLAN data between themselves.

See [Configuring Wide Area Bonjour for BGP EVPN VXLAN Layer 3 Overlay Network](#) for detailed configuration steps for Cisco Catalyst 9000 Series switches.

See [Cisco Wide Area Bonjour Application on Cisco DNA Center User Configuration Guide, 2.1.2 Release](#) for more information about how to configure Cisco DNA Center.

How to Configure Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks


This section shows how to configure Local and Wide Area Bonjour over Layer 3 overlay networks.

Configuring Local Area Bonjour in Unicast Mode over EVPN VXLAN Layer 3 Overlay Networks

Local Area Bonjour enables unicast-based mDNS service-routing function between the service provider and receiver within the local Layer 2 network boundary. The leaf switch can be directly or indirectly attached to one of:

- Wired endpoint using an intermediate Layer 2 access switch
- Wireless endpoint using a Cisco Catalyst 9800 Series WLC

Figure 87: Local Area Bonjour Unicast Mode for Wired and Wireless Local Mode Networks



Service-routing for Local Area Bonjour enables unicast-based mDNS service discovery and distribution in the access layer through the BGP EVPN leaf switch. In the distribution layer, service-routing achieves this with multiple downstream Layer 2 Ethernet switches or Cisco WLCs in Service-Peer mode. To configure Local Area Bonjour in unicast mode over an EVPN VXLAN Layer 3 overlay network for the deployment shown above, perform the following operations:

- (Optional) Configure the Layer 2 access switch in Service-Peer mode to enable the first hop mDNS gateway that directly connects it to mDNS wired and wireless endpoints.

- (Optional) Configure the Cisco Catalyst 9800 series WLC in Service-Peer mode to enable the first hop mDNS gateway that directly connects it to mDNS wired and wireless endpoints.
- Configure the Layer 3 BGP EVPN access leaf switch as an SDG agent.
- Configure the Layer 3 BGP EVPN distribution leaf switch as an SDG agent.




Note Ensure that you BGP EVPN VXLAN before you enable an mDNS gateway on a VLAN mapped to a Layer 3 overlay network. For detailed information and configuration steps, see *BGP EVPN VXLAN Configuration Guide* for the applicable release.

Configuring Layer 3 BGP EVPN Access Leaf SDG Agents

This sub-section shows a sample configuration to enable unicast-based service-routing and mDNS gateway functionality on a Layer 3 BGP EVPN access leaf switch for the topology in the figure below:

Figure 88: Layer 3 BGP EVPN Leaf Access SDG Agents



357232


Table 72: Configuring Layer 3 BGP EVPN Access Leaf SDG Agents

Configuration Step	Layer 3 Access Leaf SDG Agent in VLAN 10	Layer 3 Access Leaf SDG Agent in VLAN 20
Enable mDNS gateway and set the gateway mode on the Layer 3 Access switch.	! mdns-sd gateway mode sdg-agent !	mdns-sd gateway mode sdg-agent !
Create a unique mDNS inbound policy to permit ingress AirPrint service announcement from the service provider.	! mdns-sd service-list LOCAL-AREA-SERVICES-IN in match printer-ipp !	! mdns-sd service-list LOCAL-AREA-SERVICES-IN in match printer-ipp !
Create a unique mDNS outbound policy to permit egress AirPrint service response to the service receiver. Associate location-filter to share AirPrint service information from the grouped VLAN	! mdns-sd service-list LOCAL-AREA-SERVICES-OUT out match printer-ipp !	! mdns-sd service-list LOCAL-AREA-SERVICES-OUT out match printer-ipp !
Associate the inbound and outbound service lists to a unique service-policy.	! mdns-sd service-policy LOCAL-AREA-POLICY service-list LOCAL-AREA-SERVICES-IN service-list LOCAL-AREA-SERVICES-OUT !	! mdns-sd service-policy LOCAL-AREA-POLICY service-list LOCAL-AREA-SERVICES-IN service-list LOCAL-AREA-SERVICES-OUT !
Enable unicast-based mDNS gateway for wired users in VLAN 10 and VLAN 20 that are mapped to the Layer 3 VRF for BGP EVPN VXLAN. Configure the service-policy with advanced parameters and associate it with the VLAN.	! ! mDNS Unicast based gateway ! vlan configuration 10 mdns-sd gateway service-policy LOCAL-AREA-POLICY active-query timer 3600 !	! ! mDNS Unicast based gateway ! vlan configuration 20 mdns-sd gateway service-policy LOCAL-AREA-POLICY active-query timer 3600 !
Configure the wired users for VLAN 10 and VLAN 20 with IP address, Layer 3 VRF, and any other required parameters.	! interface Vlan 10 description BLUE VRF WIRED USER VLAN vrf forwarding BLUE_VRF ip address 10.1.1.254 255.255.255.0 no shutdown !	! interface Vlan 20 description BLUE VRF WIRED USER VLAN vrf forwarding BLUE_VRF ip address 10.1.2.254 255.255.255.0 no shutdown !

Configuring the BGP EVPN Leaf Switch and the Layer 2 Access Switch for Service-Routing

This sub-section shows a sample configuration to enable unicast-based service-routing between a Layer 2 access switch and a BGP EVPN VXLAN distribution leaf switch for the topology in the figure below:

Figure 89: Layer 2 Access Switch Service-Peer and Distribution Leaf Switch SDG Agent



357233

Table 73: Configuring the BGP EVPN Leaf Switch and the Layer 2 Access Switch for Service-Routing


Configuration Step	Layer 2 Access Switch as Service Peer	Distribution Leaf as SDG Agent
Step 1: Enable mDNS gateway and the gateway modes on the access switch (SP-1) and distribution switch (Agent-1).	! mdns-sd gateway mode service-peer !	! mdns-sd gateway mode sdg-agent !
Create a unique mDNS inbound policy to permit ingress AirPrint service announcement from the service provider on the access switch (SP-1) .	! mdns-sd service-list LOCAL-AREA-SERVICES-IN in match printer-ipp !	!

Configuration Step	Layer 2 Access Switch as Service Peer	Distribution Leaf as SDG Agent
Step 3: Create a unique mDNS outbound policy to permit egress AirPrint service response to the locally attached wired service receiver.	<pre>! mdns-sd service-list LOCAL-AREA-SERVICES-OUT out match printer-ipp !</pre>	!
Step 4: Associate the inbound and outbound service lists to a unique service-policy.	<pre>! mdns-sd service-policy LOCAL-AREA-POLICY service-list LOCAL-AREA-SERVICES-IN service-list LOCAL-AREA-SERVICES-OUT !</pre>	!
Step 5: Enable unicast-based mDNS gateway on wired users in VLAN 10 mapped to the Layer 3 VRF for BGP EVPN VXLAN. Associate the service-policy with advanced parameters. On the Service-Peer, configure the network IP address for the SDG-Agent's management VLAN and the parameters for the local source management VLAN to enable unicast service-routing over Layer 2 trunk interface.	<pre>! ! mDNS Unicast based gateway ! vlan configuration 10 mdns-sd gateway service-policy LOCAL-AREA-POLICY active-query timer 3600 sdg-agent 10.0.1.254 source-interface Vlan 4094 !</pre>	<pre>! mDNS Unicast based gateway ! vlan configuration 10 mdns-sd gateway !</pre>
Step 6: Configure the wired user in VLAN 10 with the IP address, Layer 3 VRF, and any other parameters, as required.	!	<pre>! interface Vlan 10 description BLUE VRF WIRED USER VLAN vrf forwarding BLUE_VRF ip address 10.1.1.254 255.255.255.0 no shutdown !</pre>

Configuring the Layer 2 Cisco WLC and the BGP EVPN Leaf Switch for Service-Routing

This sub-section shows a sample configuration to enable unicast-based service-routing between a Cisco Catalyst 9800 Wireless LAN Controller and a BGP EVPN VXLAN distribution leaf switch for the topology in the figure below:

Figure 90: Catalyst 9800 WLC Service-Peer and Distribution Leaf Switch SDG Agent



357234

Table 74: Configuring the Layer 2 Cisco WLC and the BGP EVPN Leaf Switch for Service-Routing

Configuration Step	WLC Service-Peer	Layer 3 Distribution Leaf Switch SDG Agent
<p>Step 1: Enable mDNS gateway and set the gateway mode.</p> <ul style="list-style-type: none"> WLC Service Peer: Service Peer mode is the default mode with mDNS gateway configuration. Layer 3 Distribution Leaf Switch: SDG Agent mode is the default mode with mDNS gateway configuration. 	<pre>! mdns-sd gateway !</pre>	<pre>! mdns-sd gateway mode sdg-agent !</pre>

Configuration Step	WLC Service-Peer	Layer 3 Distribution Leaf Switch SDG Agent
Step 2: Create a unique mDNS inbound policy to permit ingress AirPrint service announcement on the WLC in Service-Peer mode.	! <pre>mdns-sd service-list LOCAL-AREA-SERVICES-IN in match printer-ipp !</pre>	!
Step 3: Create a unique mDNS outbound policy to permit egress AirPrint service response on the WLC in Service-Peer mode.	! <pre>mdns-sd service-list LOCAL-AREA-SERVICES-OUT out match printer-ipp !</pre>	!
Step 4: Associate the inbound and outbound service lists to a unique service-policy.	! <pre>mdns-sd service-policy LOCAL-AREA-POLICY service-list LOCAL-AREA-SERVICES-IN service-list LOCAL-AREA-SERVICES-OUT !</pre>	!
Step 5: Activate unicast mDNS gateway and attach the service-policy on the WLAN and the wired VLAN. <ul style="list-style-type: none"> • WLC: Activate mDNS gateway for each WLAN Policy Profile and SSID. • Switch: Activate mDNS gateway for each VLAN. 	! <pre>wireless profile policy WLAN-PROFILE shutdown mdns-sd service-policy LOCAL-AREA-POLICY no shutdown ! wlan WLAN-PROFILE 1 blizzard shutdown mdns-sd-interface gateway no shutdown !</pre>	! <pre>! mDNS Unicast based gateway ! vlan configuration 30 mdns-sd gateway !</pre>
Step 6: (Optional) Enable service-routing on the wired Service-Peer mDNS between the local VLANs. Enable location-based wireless service on the WLC.	! <pre>mdns-sd service-policy LOCAL-AREA-POLICY location ap-location !</pre>	!
Step 7: Enable unicast service-routing between the wired and wireless Service-Peer and SDG Agent. <p>Configure the SDG Agent's IP address and the wireless management source VLAN's ID and IP address on the WLC.</p>	! <pre>mdns-sd gateway source-interface vlan 4094 sdg-agent 10.0.2.254 !</pre>	!

Configuration Step	WLC Service-Peer	Layer 3 Distribution Leaf Switch SDG Agent
Step 8: Configure the wireless user VLAN with the IP address, Layer 3 VRF, and any other parameters, as required.	!	! interface Vlan 30 description BLUE VRF WIRELESS USER VLAN vrf forwarding BLUE_VRF ip address 10.1.3.254 255.255.255.0 no shutdown !

Configuring Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks

To configure Wide Area Bonjour and connect with wired or wireless endpoints across the IP network boundary, enable unicast-mode service-routing on:

- the Layer 3 access leaf switch SDG agent
- the WLC distribution leaf switch SDG agent


Once you configure both SDG agents, you enable service-routing peering with Cisco DNA Center Wide Area Bonjour application that supports BGP EVPN VXLAN. Once you enable service-routing, global discovery and distribution is enabled and the Layer 3 leaf switch communicates with Cisco DNA Center to perform service-discovery beyond its boundary. mDNS services can be discovered and distributed among the endpoints in the same or in different Layer 3 overlay networks.

For detailed steps about how to configure Wide Area Bonjour service-routing and service policies on a Cisco DNA Center, see *Wide Area Bonjour Configuration Guidelines* in [Cisco Wide Area Bonjour Application on Cisco DNA Center User Configuration Guide, 2.1.2 Release](#).



Note We recommend that you use the Loopback interface to establish service-routing between the leaf switches and Cisco DNA Center in the underlay network.

Figure 91: Deployment Model for Wide Area Bonjour over BGP EVPN VXLAN



The following table shows how to enable unicast-based service-routing on Layer 3 overlay networks connected to wired endpoints or wireless endpoints across IP network boundary:

Table 75: Configuring SDG Agents to Configure Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks

Configuration Step	Layer 3 Access Leaf SDG Agent	WLC Dist
Step 1: Enable mDNS gateway and set the gateway mode.	! mdns-sd gateway !	! mdns-sd !
Step 2: Activate the unicast mDNS gateway on the wired VLAN and Wireless user VLAN on the respective SDG Agents.	! vlan configuration 10 mdns-sd gateway !	! vlan con mdns-sc !
Step 3: Enable Wide Area Bonjour service-routing. Configure the service export with the associated controller parameters like IP Address and source interface for stateful connection for Wide Area service-routing. Inbound policy towards the controller is not required.	! service-export mdns-sd controller DNAC-CONTROLLER-POLICY controller-address 100.0.0.1 controller-source-interface LOOPBACK 0	! service- control control control LOOPBACK !

Configuration Step	Layer 3 Access Leaf SDG Agent	WLC Distribu
Step 4: (Optional) Create a custom controller service list to permit egress AirPrint service discovery and distribution from the Catalyst Switch in SDG Agent mode.	<pre>! mdns-sd controller service-list WIDE-AREA-SERVICES match printer-ipp !</pre>	<pre>! mdns-sd cor service-lis match print !</pre>
Step 5: (Optional) Associate the controller service list to a custom controller service policy.	<pre>! mdns-sd controller service-policy WIDE-AREA-POLICY service-list WIDE-AREA-SERVICES !</pre>	<pre>! mdns-sd cor WIDE-AREA-P service-lis</pre>
Step 6: (Optional) Associate the custom controller service policy to Wide Area Service routing. Note A default controller service policy is configured if you do not create a custom controller service policy.	<pre>! service-export mdns-sd controller DNAC-CONTROLLER-POLICY controller-service-policy WIDE-AREA-POLICY !</pre>	<pre>! service-exp DNAC-CONTRC controller- WIDE-AREA-P !</pre>

Verifying Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks

This section shows how to verify Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 overlay networks.

Verifying Local Area Bonjour over an EVPN VXLAN Layer 3 Overlay Network

This section shows how to verify Local Area Bonjour over EVPN VXLAN Layer 3 overlay networks. The examples in this section show the mDNS configuration, service discovery status, and service distribution status for the policies applied on devices in Service-Peer and SDG Agent modes.

Verifying Wired Service-Peer Configuration

The examples in this section show how to verify the wired Service-Peer configuration on the Cisco Catalyst Series switch (SP-1) in Service-Peer mode. You can determine the operational status once you configure and discover AirPrint service from the local network.

This example shows a sample output for the **show mdns-sd summary vlan *vlan-id*** command on SP-1:

```
SP-1# show mdns-sd summary vlan 10
VLAN : 10
=====
mDNS Gateway : Enabled
mDNS Service Policy : LOCAL-AREA-POLICY
```

```

Active Query   : Enabled
Periodicity    : 3600 Seconds
Transport Type : IPv4
Service Instance Suffix : Not-Configured
mDNS Query Type : ALL
SDG Agent IP   : 10.0.1.254
Source Interface : Vlan4094

```

```
SP-1#
```

This example shows a sample output for the **show mdns-sd service-policy name *policy-name*** command on SP-1:

```

SP-1# show mdns-sd service-policy name LOCAL-AREA-POLICY
Service Policy Name  Service List IN Name  Service List Out Name
=====
LOCAL-AREA-POLICY   LOCAL-AREA-SERVICES-IN LOCAL-AREA-SERVICES-OUT

```

```
SP-1#
```

This example shows a sample output for the **show mdns-sd cache vlan *vlan-id*** command on SP-1:

```

SP-1# show mdns-sd cache vlan 10
NAME                                     TYPE          TTL/Remaining  Vlan-Id/If-name
Mac Address      RR Record Data
_universal._sub._ipp._tcp.local         PTR           4500/4486      V110
ac18.2651.03fe   Bldg-1-FL1-PRN._ipp._tcp.local         PTR           4500/4486      V110
_ac18.2651.03fe   Bldg-1-FL1-PRN._ipp._tcp.local         SRV           4500/4486      V110
Bldg-1-FL1-PRN._ipp._tcp.local         0 0 631 Bldg-1-FL1-PRN.local                    A             4500/4486      V110
ac18.2651.03fe   10.1.1.1                                     AAAA          4500/4486      V110
ac18.2651.03fe   2001:10:153:1:79:A40C:6BEE:AEEC          TXT           4500/4486      V110
Bldg-1-FL1-PRN._ipp._tcp.local         (451)'txtvers=1''priority=30''ty=EPSON WF-3620
Series''usb_MFG=EPSON''usb_MDL=W~'~

```

```
SP-1#
```

This example shows a sample output for the **show mdns-sd statistics vlan *vlan-id*** command on SP-1:

```

SP-1# show mdns-sd statistics vlan 10
mDNS Statistics

V110:
mDNS packets sent : 612
  IPv4 sent : 612
    IPv4 advertisements sent : 0
    IPv4 queries sent : 612
  IPv6 sent : 0
    IPv6 advertisements sent : 0
    IPv6 queries sent : 0
Unicast sent : 0
mDNS packets rate limited : 0
mDNS packets received : 42
  advertisements received : 28
  queries received : 14
  IPv4 received : 42

```

```

IPv4 advertisements received : 28
IPv4 queries received : 14
IPv6 received : 0
IPv6 advertisements received : 0
IPv6 queries received : 0
mDNS packets dropped : 0
=====
Query Type                : Count
=====
PTR                        : 12
SRV                        : 0
A                          : 0
AAAA                      : 0
TXT                        : 0
ANY                        : 3
=====
PTR Name                   Advertisement      Query
=====
_ipp._tcp.local            9              4
SP-1#

```

Verifying Wired SDG Agent Configuration and Service-Routing Status

The examples in this section show how to verify the mDNS configuration and service-routing on Wired SDG Agent (SDG-1). SDG-1 is locally connected to a Layer 2 access switch (SP-1) in Service-Peer mode and centrally paired with a Cisco DNA Center for Wide Area Bonjour service-routing.

This example shows a sample output for the **show mdns-sd summary vlan *vlan-id*** command on SDG-1:

```

SP-1# show mdns-sd summary vlan 10
VLAN : 10
=====
mDNS Gateway      : Enabled
mDNS Service Policy : LOCAL-AREA-POLICY
Active Query      : Disabled
Transport Type    : IPv4
Service Instance Suffix : Not-Configured
mDNS Query Type   : ALL
SDG Agent IP      : Not-Configured
Source Interface   : Not-Configured
SDG-1#

```

This example shows a sample output for the **show mdns-sd cache vlan *vlan-id*** command on SDG-1:

```

SP-1# show mdns-sd cache vlan 10
NAME                               TYPE          TTL/Remaining  Vlan-Id/If-name
-----
Mac Address      RR Record Data
_universal._sub._ipp._tcp.local    PTR           4500/4500      V110
ac18.2651.03fe    Bldg-1-FL1-PRN._ipp._tcp.local
_ipp._tcp.local    PTR           4500/4500      V110
ac18.2651.03fe    Bldg-1-FL1-PRN._ipp._tcp.local
Bldg-1-FL1-PRN._ipp._tcp.local    SRV           4500/4500      V110
ac18.2651.03fe    0 0 631 Bldg-1-FL1-PRN.local
Bldg-1-FL1-PRN.local              A             4500/4500      V110
ac18.2651.03fe    10.1.1.1
Bldg-1-FL1-PRN.local              AAAA          4500/4500      V110
ac18.2651.03fe    2001:10:153:1:79:A40C:6BEE:AEEC

```

```
Bldg-1-FL1-PRN._ipp._tcp.local      TXT      4500/4500      V110
ac18.2651.03fe      (451)'txtvers=1''priority=30''ty=EPSON WF-3620
Series''usb_MFG=EPSON''usb_MDL=W~'~
```

```
SDG-1#
```

This example shows a sample output for the **show mdns-sd sp-sdg statistics** command on SDG-1:

```
SP-1# show mdns-sd sp-sdg statistics
                                One min, 5 mins, 1 hour
Average Input rate (pps)      :      0,      0,      0
Average Output rate (pps)    :      0,      0,      0
Messages received:
  Query                       : 15796
  ANY query                   : 0
  Advertisements              : 28
  Advertisement Withdraw      : 0
  Interface down              : 0
  Vlan down                   : 0
  Service-peer ID change      : 0
  Service-peer cache clear    : 12
  Resync response             : 6
Messages sent:
  Query response              : 5975
  ANY Query response          : 0
  Cache-sync                  : 61
  Get service-instance        : 0
```

```
SDG-1#
```

This example shows a sample output for the **show mdns-sd controller detail** command on SDG-1:

```
SP-1# show mdns-sd controller detail

Controller : DNAC-Policy
IP : 100.0.0.1, Dest Port : 9991, Src Port : 42446, State : UP
Source Interface : Loopback0, MD5 Disabled
Hello Timer 30 sec, Dead Timer 120 sec, Next Hello 00:00:24
Uptime 2d05h (17:02:37 UTC Jan 15 2021)
Service Buffer : Enabled

Service Announcement :
Filter : DNAC-CONTROLLER-POLICY
Count 50, Delay Timer 30 sec, Pending Announcement 0, Pending Withdraw 0
Total Export Count 56, Next Export in 00:00:24

Service Query :
Query Suppression Enabled
Query Count 50, Query Delay Timer 15 sec, Pending 0
Total Query Count 15791, Next Query in 00:00:09
```

```
SDG-1#
```

Verifying Wireless Service-Peer and Service Routing Status

The examples in this section show how to verify the wireless Service-Peer configuration on the Cisco Catalyst 9800 WLC (SP-2) in Service-Peer mode. You can determine the operational status once you configure and discover AirPrint service from the remote network.

This example shows a sample output for the **show mdns-sd summary** command on SP-2:

```

SP-2# show mdns-sd summary
mDNS Gateway: Enabled
Mode: Service Peer
Service Announcement Periodicity(in seconds): 30
Service Announcement Count: 50
Service Query Periodicity(in seconds): 15
Service Query Count: 50
Active Response Timer (in seconds): Disabled
ANY Query Forward: Disabled
SDG Agent IP: 10.0.2.254
Source Interface: Vlan4094
Active Query Periodicity (in minutes): 15
Transport Type: IPv4
mDNS AP service policy: default-mdns-service-policy

```

```
SP-2#
```

This example shows a sample output for the **show wireless profile policy detailed wireless-profile-name** command on SP-2:

```
SP-2# show wireless profile policy detailed WLAN-PROFILE | sec mDNS
```

```

mDNS Gateway
  mDNS Service Policy name           : LOCAL-AREA-POLICY

```

```
SP-2#
```

This example shows a sample output for the **show mdns-sd statistics wlan-id wlan-id-value** command on SP-2:

```

SP-2# show mdns-sd statistics wlan-id 1
mDNS Packet Statistics
-----
mDNS stats last reset time: 01/10/21 21:38:19
mDNS packets sent: 4592
  IPv4 sent: 4592
    IPv4 advertisements sent: 4592
    IPv4 queries sent: 0
  IPv6 sent: 0
    IPv6 advertisements sent: 0
    IPv6 queries sent: 0
  Multicast sent: 0
    IPv4 sent: 0
    IPv6 sent: 0
mDNS packets received: 297
  advertisements received: 80
  queries received: 217
  IPv4 received: 297
    IPv4 advertisements received: 80
    IPv4 queries received: 217
  IPv6 received: 0
    IPv6 advertisements received: 0
    IPv6 queries received: 0
mDNS packets dropped: 297
Query Type Statistics
  PTR queries received: 1720
  SRV queries received: 8
  A queries received: 8
  AAAA queries received: 8
  TXT queries received: 97
  ANY queries received: 153
  OTHER queries received: 0

```



```
SP-2#
```

This example shows a sample output for the **show mdns-sd sp-sdg statistics** command on SP-2:

```
SP-2# show mdns-sd sp-sdg statistics
mDNS SP Statistics
last reset time: 01/10/21 21:37:36

Messages sent:
  Query                : 12675
  ANY query            : 0
  Advertisements       : 24
  Advertisement Withdraw : 0
  Service-peer ID change : 0
  Service-peer cache clear : 7
  Resync response      : 5
Messages received:
  Query response       : 4619
  ANY Query response   : 0
  Cache-sync           : 48
  Get service-instance : 0
```

```
SP-2#
```

This example shows a sample output for the **show mdns-sd query-db** command on SP-2:

```
SP-2# show mdns-sd query-db
MDNS QUERY DB

Client MAC: 4c32.7593.e3af
Vlan ID: 30
Wlan ID: 1
Location Group ID: 0
PTR Name(s):
  _ipp._tcp.local
```

```
SP-2#
```

Verifying Wide Area Bonjour over EVPN VXLAN Layer 3 Overlay Networks

This section shows how to verify Wide Area Bonjour over EVPN VXLAN Layer 3 overlay networks.

The examples in this section show the mDNS configuration and service-routing on a wireless SDG-Agent (SDG-2). SDG-2 is locally connected to a Catalyst 9800 series WLC (SP-2) in Service-Peer mode and centrally paired with a Cisco DNA Center for Wide Area Bonjour service-routing.

This example shows a sample output for the **show mdns-sd summary vlan *vlan-id*** command on SDG-2.

```
SDG-2# show mdns-sd summary vlan 30
VLAN : 30
=====
mDNS Gateway          : Enabled
mDNS Service Policy   : LOCAL-AREA-POLICY
Active Query          : Disabled
Transport Type        : IPv4
Service Instance Suffix : Not-Configured
mDNS Query Type       : ALL
```

```
SDG Agent IP      : Not-Configured
Source Interface  : Not-Configured
```

```
SDG-2#
```

This example shows a sample output for the **show mdns-sd sp-sdg statistics** command on SDG-2.

```
SDG-2# show mdns-sd sp-sdg statistics
                                One min, 5 mins, 1 hour
Average Input rate (pps) :      0,      0,      0
Average Output rate (pps) :      0,      0,      0
Messages received:
  Query                       : 12191
  ANY query                    : 0
  Advertisements               : 0
  Advertisement Withdraw      : 0
  Interface down               : 0
  Vlan down                    : 0
  Service-peer ID change      : 0
  Service-peer cache clear    : 18
  Resync response              : 10
Messages sent:
  Query response               : 1975
  ANY Query response           : 0
  Cache-sync                   : 19
  Get service-instance         : 0
```

```
SDG-2#
```

This example shows a sample output for the **show mdns-sd controller detail** command on SDG-2.

```
SDG-2# show mdns-sd controller detail
Controller : DNAC-Policy
IP : 100.0.0.1, Dest Port : 9991, Src Port : 42931, State : UP
Source Interface : Loopback0, MD5 Disabled
Hello Timer 30 sec, Dead Timer 120 sec, Next Hello 00:00:19
Uptime 2d05h (17:10:18 UTC Jan 15 2021)
Service Buffer : Enabled

Service Announcement :
Filter : DNAC-CONTROLLER-POLICY
Count 50, Delay Timer 30 sec, Pending Announcement 0, Pending Withdraw 0
Total Export Count 0, Next Export in 00:00:19

Service Query :
Query Suppression Enabled
Query Count 50, Query Delay Timer 15 sec, Pending 0
Total Query Count 17093, Next Query in 00:00:19
```

```
SDG-2#
```

Verifying Cisco DNA Center Configuration and Service-Routing Status

The Cisco Wide Area Bonjour application supports multilevel assurance capabilities for day-to-day operations. The service routing, instance monitoring, management, and troubleshooting in the Cisco Wide Area Bonjour application are divided into three major categories. Each category serves a unique function to manage and troubleshoot Wide Area Bonjour service routing for day-2 operation. The Monitor function comprises the following categories:

- **Dashboard:** The landing page of Cisco Wide Area Bonjour application provides an overview of key statistics in multiple formats. It allows you to quickly determine the service-routing health across the network. The network administrator can monitor the operational status of service-routing using:
- **Sub-Domain 360°:** Network administrators can collect statistics and status counts in the 360° view. The left-panel monitoring and configuration bar opens automatically upon clicking a subdomain. You can verify configured policies and discovered service instances on a per-subdomain basis of the Configuration section.
- **Monitor:** A comprehensive three-tier monitoring and troubleshooting function of the Cisco Wide Area Bonjour application for various day-2 operations. Network administrators use the detail view of the SDG Agent, service instance, and advanced troubleshooting capabilities to manage and troubleshoot the Wide Area Bonjour domain with a single pane of glass in Cisco DNA Center.

For detailed information about assurance capabilities and monitoring of operation details, see [Monitor the Cisco Wide Area Bonjour Application](#) module of Cisco Wide Area Bonjour on Cisco DNA Center User Guide, Release 2.1.2.

Additional References for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks

Table 76: Additional References for Cisco DNA Service for Bonjour over EVPN VXLAN Layer 3 Overlay Networks

Related Topic	Document Title
DNA Service for Bonjour Deployment on Cisco Catalyst 9300 Series Switches	Cisco DNA Service for Bonjour Configuration Guide (Catalyst 9300 Switches)
DNA Service for Bonjour Deployment on Cisco Catalyst 9800 Wireless LAN Controllers	Cisco Catalyst 9800 Series Wireless Controller Software Configuration Guide
Cisco Wide Area Bonjour Application on Cisco DNA Center User Guide	Cisco Wide Area Bonjour Application on Cisco DNA Center User Guide, Release 2.1.2



CHAPTER 22

Configuring VRF-Aware Local Area Bonjour Services

Beginning from Cisco IOS XE Bengaluru 17.4.1, Cisco Catalyst 9000 Series switches supports Virtual Routing and Forwarding-Aware (VRF-Aware) services in Local Area Bonjour domain. VRF-Aware Local Area Bonjour services provide boundary-based service discovery for Layer 3 segmented IPv4 and IPv6 network and support policy-based (secure) routing services for Wired and Wireless networks. VRF-Aware Local Area Bonjour service is supported on enterprise-grade, traditional, and next-generation fabric-based deployment models as described in [Cisco DNA Service for Bonjour Solution Overview](#).

- [Prerequisites for VRF-Aware Local Area Bonjour Services, on page 911](#)
- [Restrictions for VRF-Aware Local Area Bonjour Services, on page 912](#)
- [Information about VRF-Aware Local Area Bonjour Services, on page 912](#)
- [Understanding VRF-Aware Wide Area Bonjour Services, on page 914](#)
- [Understanding VRF-Aware Service on Multilayered Wired and Wireless Networks, on page 915](#)
- [How to configure Intra-Virtual Network Proxy Service on Local Area Bonjour Domain, on page 917](#)
- [How to configure Inter-Virtual Network Proxy Service on Local Area Bonjour Domain, on page 918](#)
- [Verifying VRF-Aware Local Area Bonjour Services, on page 921](#)

Prerequisites for VRF-Aware Local Area Bonjour Services

- You must understand the mDNS service segmentation capabilities to implement, manage, and troubleshoot the proxy service in Local Area Bonjour domain.
- Ensure that the Cisco Catalyst 9000 Series switch is configured in SDG-Agent mode. VRF-Aware Local Area Bonjour service is supported on first-hop IP gateway of switches configured in SDG-Agent mode in Wired and Wireless networks.
- Ensure that the software version installed on the Cisco Catalyst 9000 Series switch is Cisco IOS XE Bengaluru 17.4.1 or higher.
- Ensure that all required IP VRF with IPv4 or IPv6 address-family configurations is completed. These configurations are required to activate VRF on the switch configured in SDG-Agent mode.
- Ensure that the IP VRF configured to a local SVI interface supports IP gateway so that the mDNS Wired and Wireless endpoint can be attached directly or remotely.

- To activate mDNS gateway in Unicast mode for a VLAN, ensure that the mDNS gateway and service policy is configured after enabling the VLAN using the **vlan configuration id** command.
- Ensure that all configurations for IPv4 or IPv6-based data routing and forwarding both within the same VRF or different VRFs are complete including network requirements such as stateful firewall configuration, route-leaking configuration and so on.
- Ensure that all the prerequisites described in *Configuring Local Area Bonjour in Unicast Mode for LAN Networks* module are completed.

Restrictions for VRF-Aware Local Area Bonjour Services

- VRF-Aware Local Area Bonjour service is not supported on a Layer 2 Cisco Catalyst 9000 Series switch or a Layer 2 Cisco Catalyst 9800 WLC in Service-Peer mode.
- VRF-Aware Local Area Bonjour services are configured to provide mDNS service discovery information between Layer 3 segments within the same or different IP VRF, or share services from non-VRF enabled networks only. Any additional IP routing and data forwarding configurations are beyond the scope of this implementation.

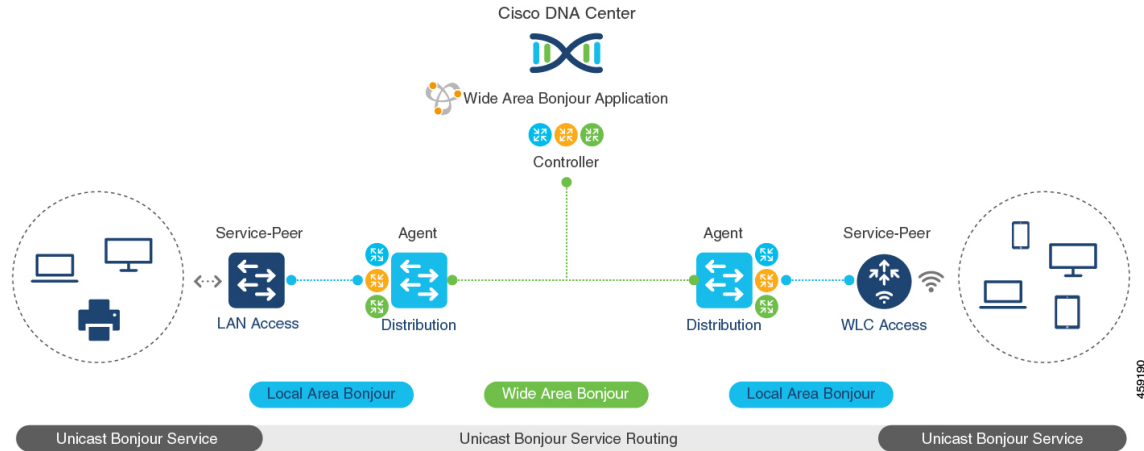
Information about VRF-Aware Local Area Bonjour Services

The Cisco DNA Service for Bonjour solution provides end-to-end service-routing for enterprise-grade Wired and Wireless networks. The enterprise network builds secure and segmented networks that protect IT-managed infrastructure and shares services and resources among trusted and untrusted user group. The physical infrastructure can be logically virtualized into a private networking space that supports secure communication services within closed user groups and conditionally extends boundary services based on business and technical demands.

VRF-Aware Local Area Bonjour gateway services allow to dynamically discover and distribute mDNS services on the same VRF segmented Layer 3 overlay networks based on policy. You can also build an Extranet network using the mDNS location-filter policy that supports proxy services among multiple logical VRF or a global IP routing domain on a local system. The Layer 3 VRF segmented networks can also be configured to route in overlay using any next-generation overlay networks such as Cisco SD-Access, BGP EVPN VXLAN or classic technologies such as Multi-VRF, MPLS.

[Figure](#) illustrates the Cisco DNA Service for Bonjour solution configured with VRF-Aware services for enterprise-grade Wired and Wireless networks.

Figure 92: Cisco DNA Service for Bonjour with VRF-Aware Services




459190

Gateway Modes for VRF-Aware Bonjour Services

VRF is a Layer 3 specific virtual routing function and therefore it is implemented on Layer 3 Ethernet switches with first-hop IP gateways that can directly or remotely attach mDNS endpoints.

Figure illustrates the Cisco Catalyst 9000 Series switch in SDG-Agent mode supporting VRF-Aware services in Layer 3 access mode and in multi-layer network deployment mode. In multi-layer network deployment mode, the gateway to the distribution layer provides a Layer 2 or Layer 3 boundary to a downstream Layer 2 Cisco Catalyst 9000 Series switch and Cisco Catalyst 9800 WLC for local proxy service with local VLANs.

Figure 93: Gateway Modes for VRF-Aware Services



459191

- **VRF-Aware Routed Access:** The Cisco Catalyst 9000 Series switch can be deployed as an IP gateway for directly attached Wired or Wireless mDNS endpoints. The Cisco Wireless SSID can be configured as fabric-enabled or as FlexConnect with local switching that provides local termination point to a first-hop Ethernet switch that supports Layer 3 overlay networks such as Cisco SD-Access or BGP EVPN based-fabric networks. A Cisco Catalyst 9000 Series switch configured in SDG-Agent mode provides unicast-based mDNS gateway services to directly attached Wired and Wireless endpoints within the same or different virtual routing network space or a default global IP network.
- **VRF-Aware Multilayer:** The Cisco Catalyst 9000 Series Switch can be deployed as an IP gateway for remotely attached Wired or Wireless mDNS endpoints through an intermediate Layer 2 Cisco Catalyst 9000 Series switch or Cisco Catalyst 9800 Series WLC. A Cisco Catalyst 9000 Series switch, configured in SDG-Agent mode and in the distribution layer, provides VRF-Aware mDNS gateway services, while the Layer 2 Ethernet switch and Cisco WLC in Unicast mode provides local proxy services to directly attached Wired and Wireless endpoints within the same or different VLAN.


Understanding VRF-Aware Wide Area Bonjour Services

The VRF-Aware service discovery and distribution can be implemented across multiple switches in SDG-Agent mode on an IP, MPLS, or VXLAN-enabled network with Wide Area Bonjour. The Cisco DNA-Center Wide Area Bonjour application supports granular and policy-based routing services that allow discovery and distribution of mDNS services dynamically over overlay networks. You can build a global policy combining one or more source and receiver SDG-Agents that allow distributing or advertising services from a specific IPv4 or even an IPv6 network mapped to the VRF.

The network wide and distributed switches in SDG-Agent mode transport locally discovered or requested mDNS service information over lightweight unicast routing services to a centralized Cisco DNA-Center controller in an underlay IPv4 network. These switches must be configured with a unified service-export policy for local networks mapped to one or more VRFs or to a global IP routing domain.

[Figure](#) illustrates VRF-Aware Wide Area Bonjour services for IP, MPLS, or VXLAN enabled overlay networks.

Figure 94: VRF-Aware Wide Area Bonjour Services



The *Configuring Wide Area Bonjour* module lists the configuration procedures in detail.

Understanding VRF-Aware Service on Multilayered Wired and Wireless Networks


The Cisco Catalyst 9000 Series switches support VRF-Aware service for multilayered Wired and central-switching Wireless-enabled networks. The Layer 2 or Layer 3 network boundary to the Cisco Catalyst 9000 Series switches is extended at the distribution layer with an intermediate Layer 2 Cisco Catalyst 9000 Series switch or Cisco Catalyst 9800 Series WLC and directly attached to the Wired and central-switching Wireless endpoints. As the IP gateway shifts, the Cisco Catalyst 9000 Series switches in the distribution layer must be configured in SDG-Agent mode and the downstream Layer 2 switch and WLC network devices must be configured in Service-Peer mode to support mDNS proxy services to locally attached endpoints.

The VRF-Aware service configured on a switch, in SDG-Agent mode and in the distribution layer, follows configuration and operation guidelines for Wired and central-switching Wireless as described in [Understanding VRF-Aware Wide Area Bonjour Services, on page 914](#). The Layer 2 switch and WLC network devices remains transparent to VRF-Aware services and continues to provide local proxy services to locally attached users in the same or different VLANs.

The VRF-Aware service discovery and distribution can be implemented across multiple switches in SDG-Agent mode on an IP, MPLS, or VXLAN-enabled network with Wide Area Bonjour. The Cisco DNA-Center Wide Area Bonjour application supports granular and policy-based routing services that allow discovery and distribution of mDNS services dynamically for overlay networks. You can build a global policy combining one or more source and receiver SDG-Agent that allow distributing or advertising services from a specific IPv4 or even an IPv6 network mapped to the VRF.

[Figure](#) illustrates end-to-end VRF-Aware on multilayered Wired and Wireless networks across Wide Area Bonjour domain with Cisco DNA-Center.

Figure 95: VRF-Aware on Multilayered Wired and Wireless Network



How to configure Intra-Virtual Network Proxy Service on Local Area Bonjour Domain


Intra-Virtual Network (Intra-VN) Proxy Service is a policy-based VRF-Aware service discovery and distribution implemented on the IP VRF of a switch in SDG-Agent mode connected to multiple IP networks.

Beginning from Cisco IOS XE Bengaluru 17.4.1, the Cisco Catalyst 9000 Series switches support mDNS gateway service as the default on each VRF. You must build a mDNS service policy that implicitly allows required mDNS service types and mapping services to endpoint facing VLANs. The Cisco Catalyst 9000 Series switch can automatically discover VRF associations to a VLAN interface without additional configurations.

The Cisco Catalyst 9000 Series switch in SDG-Agent mode dynamically discovers mDNS services from a local network and automatically builds VRF-aware service information. To enable Layer 3 segmented proxy service by default, the SDG-Agent provides limited mDNS service proxy response to endpoints in other VLANs mapped with the same VRF.

Figure illustrates VRF-Aware enabled on an Intra-VN proxy service.

Figure 96: Intra-VN Service Proxy




How to configure Inter-Virtual Network Proxy Service on Local Area Bonjour Domain

Inter-VN Proxy Service is a policy-based VRF-Aware service discovery and distribution implemented on multiple IP VRFs or on a global IP routing domain of a switch in SDG-Agent mode connected to multiple IP networks.

Beginning from Cisco IOS XE Bengaluru 17.4.1, the Cisco Catalyst 9000 Series switches support mDNS service discovery and distribution between IP VRFs or on a global routing domain based on the configured mDNS location-filter policy. The existing location-filter configuration on an SDG-Agent permits mDNS service information between configured VLANs and records discovery and distribution on the mapping table. Although configuring inter-VN provides Extranet mDNS proxy services between Wired and Wireless networks, additional methods such as stateful firewall, route-leaking and so on must also be configured to handle the data transfer between Inter-VN or VRF to global IP routing.

Figure shows Inter-VN proxy service for Extranet network.

Figure 97: Inter-VN Proxy Service



485195

Configuring Inter-Virtual Network Location-Filter

To enable the local service proxy on the switch to discover mDNS services between local VLANs, 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 ID Example: Device(config)# vlan 101 Device(config-vlan)# name BLUE-VRF Device(config)# vlan 201 Device(config-vlan)# name YELLOW-VRF Device(config)# vlan 301 Device(config-vlan)# name GREEN-VRF	Configures a VLAN ID in local database for overlay mDNS endpoints.

	Command or Action	Purpose
Step 4	mdns-sd location-filter <i>location-filter-name</i> Example: <pre>Device(config)# mdns-sd location-filter INTER-VN-LOCAL-PROXY</pre>	Configures a unique location-filter.
Step 5	match location {all default ID} vlan [ID] Example: <pre>Device(config-mdns-loc-filter)# match location-group default vlan 101 Device(config-mdns-loc-filter)# match location-group default vlan 201 Device(config-mdns-loc-filter)# match location-group default vlan 301</pre>	Configures the match criteria that mutually distribute permitted services between grouped VLANs.
Step 6	mdns-sd service-list <i>service-list-name</i> {in out} Example: <pre>Device(config)# mdns-sd service-list BLUE-VRF-LIST-OUT out</pre>	Configures mDNS service-list to classify one or more service-types. Unique service-list is required to process incoming mDNS message and the outbound response to the requesting end points.
Step 7	match <i>service-destination-name</i> [message-type {any announcement query}] Example: <pre>Device(config)# mdns-sd service-list BLUE-VRF-LIST-OUT out Device(config-mdns-sl-out)# match APPLE-TV location-filter LOCAL-PROXY</pre>	<p>Associates the location-filter to one or more service-types to enable local proxy between local VLANs. For example, the Apple-TV learned from the YELLOW-VRF VLAN 201 and the GREEN-VRF VLAN 301 will be distributed to the receiver in the BLUE-VRF VLAN 101.</p> <p>The service-list contains implicit deny at the end.</p> <p>The message-type for outbound service-list is not required.</p>
Step 8	mdns-sd service-policy <i>service-policy-name</i> Example: <pre>Device(config)# mdns-sd service-policy BLUE-VRF-POLICY</pre>	Creates a unique mDNS service-policy.
Step 9	service-list <i>service-list-name</i> {in out} Example: <pre>Device(config)# mdns-sd service-policy BLUE-VRF-POLICY Device(config-mdns-ser-policy)# service-list BLUE-VRF-LIST-OUT out</pre>	Configures an mDNS service policy to associate with the service-list for each direction.
Step 10	vlan configuration <i>ID</i> Example: <pre>Device(config)# vlan configuration 101-103</pre>	<p>Enables VLAN configuration for advanced service parameters.</p> <p>One or more VLANs can be created for the same settings. For example, the VLAN</p>

	Command or Action	Purpose
		configuration range 101-110 or 200 allows to configure consecutive and nonconsecutive VLAN IDs.
Step 11	mdns-sd gateway Example: Device (config-vlan) # mdns-sd gateway	Enables the mDNS gateway on the specified VLAN IDs.
Step 12	service-policy BLUE-VRF-POLICY Example: Device (config-vlan-mdns) # service-policy BLUE-VRF-POLICY	Associates an mDNS service-policy with the specified VLAN IDs.
Step 13	end Example: Device (config-vlan-mdns) # end	Returns to privileged EXEC mode.

Verifying VRF-Aware Local Area Bonjour Services

The dynamically discovered VRF-Aware service information can be verified on Cisco Catalyst 9000 Series switch in SDG-Agent mode by including the **vrf** keyword on the existing **show mdns-sd** command. You can verify each VRF-service record information based on the unique VRF name.

The following is an example of the command that displays the dynamically discovered mDNS service records in the BLUE-VRF:

```
Device# show mdns-sd cache vrf BLUE-VRF
```

```

                                     mDNS CACHE
-----
[<NAME>]                               [<TYPE>]      [<TTL>/Remaining] [Vlan-Id/If-name]
[Mac Address]           [<RR Record Data>]

RTP-ATV-1._device-info._tcp.local      TXT           4500/4495         511
a018.28f2.9889          (13) 'model=J33iAP'
RTP-ATV-1._airplay._tcp.local          PTR           4500/4495         511
a018.28f2.9889          RTP-ATV-1._airplay._tcp.local
RTP-ATV-1._raop._tcp.local             PTR           4500/4495         511
a018.28f2.9889          A01828F29889@RTP-ATV-1._raop._tcp.local
RTP-ATV-1._airplay._tcp.local          SRV           4500/4495         511
a018.28f2.9889          0             0             7000             RTP-ATV-3.local
A01828F29889@RTP-ATV-1._raop._tcp.local SRV           4500/4495         511
a018.28f2.9889          0             0             7000             RTP-ATV-3.local
RTP-ATV-1.local                AAAA          4500/4495         511
a018.28f2.9889          2001:10:153:2:C2F:9445:7062:5C3C
RTP-ATV-1.local                A             4500/4495         511
a018.28f2.9889          10.155.1.17
RTP-ATV-1._airplay._tcp.local          TXT           4500/4495         511
a018.28f2.9889
(208) 'deviceid=A0:18:28:F2:98:89' 'features=0x5A7FFFFF7,0x1E' 'flags=0x44' 'model=~'~
A01828F29889@RTP-ATV-1._raop._tcp.local TXT           4500/4495         511

```

```
a018.28f2.9889
(177) 'cn=0,1,2,3'da=true'et=0,3,5'ft=0x5A7FFFF7,0x1E'md=0,1,2'am=AppleTV3,2'~
```

Use the following commands in privileged EXEC mode on a Cisco Catalyst 9000 Series switch configured in SDG-Agent mode to verify various Local Area Bonjour domain mDNS parameters such as service configuration, cache records, statistics, and so on.

Table 77: Commands to Verify VRF-Aware Services

Command	Purpose
show mdns-sd cache {all interface mac name service-peer static type vlan vrf}	<p>Displays all available mDNS cache record that supports multiple variables and provides granular source details. The following variables are available:</p> <ul style="list-style-type: none"> • all: Displays all available cache records discovered from multiple source connections of a system. • interface: Displays the available cache records discovered from a specified Layer 3 interface. • mac: Displays the available cache records discovered from the specified MAC address. • name: Displays the available cache records based on service provider announced name. • service-peer: Displays available cache records discovered from the specified Layer 2 Service-Peer. • static: Displays the locally configured static mDNS cache entries. • type: Displays the available cache records based on the specific mDNS record type (PTR, SRV, TXT, A, or AAAA). • vlan: Displays the available cache records discovered from the specified Layer 2 VLAN ID in unicast mode. • vrf: Displays each VRF available cache records based on the specific mDNS record type (PTR, SRV, TXT, A, or AAAA).
show mdns-sd service-definition {name type}	<p>Displays the built-in and user-defined custom service definitions and provides the mapping from service name to mDNS PTR records.</p> <p>The service-definition can be filtered by name or type.</p>

Command	Purpose
show mdns-sd service-list { direction name }	<p>Displays the configured inbound or outbound service-list that classifies matching service types for a service policy.</p> <p>The service lists can be filtered by name or specific direction.</p>
show mdns-sd service-policy { interface name }	<p>Displays the list of mDNS service policies mapped with inbound or outbound service-lists.</p> <p>The service policies list can be filtered by the associated specified interface or by name.</p>
show mdns-sd statistics { all cache debug interface service-list service-policy services vlan }	<p>Displays the detailed mDNS statistics processed bidirectionally by the system on each mDNS-gateway-enabled VLAN, when mDNS is configured in unicast mode.</p> <p>The keywords for the mDNS statistics provide a detail view on the interface, policy, service-list, and services.</p>
show mdns-sd summary { interface vlan }	<p>Displays the brief information about mDNS gateway and the key configuration status on all VLANs and interfaces of the system.</p>



PART II

Feature History for BGP EVPN VXLAN

- [Feature History for BGP EVPN VXLAN, on page 927](#)



CHAPTER 23

Feature History for BGP EVPN VXLAN

- [Feature History for BGP EVPN VXLAN, on page 927](#)

Feature History for BGP EVPN VXLAN

This table provides release and related information for the features explained in this module.

These features are available in all the releases subsequent to the one they were introduced in, unless noted otherwise.

Release	Feature	Feature Information
Cisco IOS XE Fuji 16.9.1	EVPN VXLAN Overlay Network for IPv4 Bridged Traffic	EVPN VXLAN overlay network for IPv4 bridged traffic is a Layer 2 overlay network that allows host devices within the same subnet to send IPv4 bridged traffic to each other using a Layer 2 virtual network instance (VNI).
	EVPN VXLAN Overlay Network for IPv4 Routed Traffic	EVPN VXLAN overlay network for IPv4 routed traffic is a Layer 3 overlay network that allows host devices in different Layer 2 networks to send IPv4 routed traffic to each other using a Layer 3 VNI and an IP VRF.
	Layer 2 Broadcast, Unknown Unicast, and Multicast (BUM) Traffic Forwarding using Underlay Multicast	Multi-destination Layer 2 broadcast, unknown unicast, and multicast (BUM) traffic in an EVPN VXLAN network is replicated through a multicast group in the underlay network and forwarded to all the endpoints of the network.
	Leaf Functionality	A leaf switch sits on the edge of a BGP EVPN VXLAN fabric and is connected to the host or access devices. It functions as a virtual tunnel end point (VTEP) and performs encapsulation and decapsulation.
	EVPN VXLAN Integrated Routing and Bridging	EVPN VXLAN integrated and bridging (IRB) allows the VTEPs in a VXLAN network to forward both Layer 2 or bridged traffic and Layer 3 or routed traffic. It is implemented as symmetric and asymmetric IRB.
	EVPN VXLAN Distributed Anycast Gateway	EVPN VXLAN distributed anycast gateway is a default gateway addressing mechanism that enables the use of the same gateway IP address across all the leaf switches that are part of a VXLAN network. Support was introduced with manual MAC address configuration on the Layer 2 VNI VLAN's switch virtual interface (SVI) on all VTEPs as the only method to enable the feature.
	DHCP Relay for IPv4 Traffic in BGP EVPN VXLAN Fabric	The VTEP in a BGP EVPN VXLAN fabric is configured as a DHCP relay agent to provide DHCP relay services for IPv4 traffic in a multi-tenant VXLAN environment.

Release	Feature	Feature Information
Cisco IOS XE Gibraltar 16.11.1	EVPN VXLAN Overlay Network for IPv6 Bridged Traffic	EVPN VXLAN overlay network for IPv6 bridged traffic is a Layer 2 overlay network that allows host devices within the same subnet to send IPv6 bridged traffic to each other using a Layer 2 VNI.
	EVPN VXLAN Overlay Network for IPv6 Routed Traffic	EVPN VXLAN overlay network for IPv6 routed traffic is a Layer 3 overlay network that allows host devices in different Layer 2 networks to send IPv6 routed traffic to each other using a Layer 3 VNI and an IP VRF.
	Layer 2 Broadcast, Unknown Unicast, and Multicast (BUM) Traffic Forwarding using Ingress Replication	Ingress replication is a unicast approach to handle multi-destination Layer 2 BUM traffic in an EVPN VXLAN network. It involves an ingress device replicating every incoming BUM packet and sending them as a separate unicast to the remote egress devices.

Release	Feature	Feature Information
Cisco IOS XE Gibraltar 16.12.1	MAC Aliasing for EVPN VXLAN Distributed Anycast Gateway	<p>MAC aliasing allows the leaf switches in an EVPN VXLAN network to advertise the MAC addresses of their Layer 2 VLAN's SVI as the gateway MAC address to all the other leaf switches in the network.</p> <p>MAC aliasing removes the need to explicitly configure the same MAC address on the Layer 2 VNI VLAN's SVI on all VTEPs in order to enable distributed anycast gateway.</p>
	Border Leaf Functionality	A border leaf switch is a leaf switch in a BGP EVPN VXLAN fabric that enables external connectivity with other Layer 2 and Layer 3 networks by acting as the connecting node between the two networks.
	Autonomous System Number Rewrite	The rewrite-evpn-rt-asn command was introduced to enable the rewrite of the autonomous system number (ASN) of the EVPN route target that originates from the current autonomous system with the ASN of the target eBGP EVPN peer.
	VRF-Lite Border Leaf Handoff	VRF-Lite border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with a VRF-Lite network through a border leaf switch.
	MPLS Layer 3 VPN Border Leaf Handoff	MPLS Layer 3 VPN border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with an MPLS Layer 3 VPN network through a border leaf switch.
	IEEE 802.1Q Border Leaf Handoff	IEEE 802.1Q border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an IEEE 802.1Q network through a border leaf switch.
	Access Border Leaf Handoff	Access border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an Access network through a border leaf switch.
	VPLS over MPLS Border Leaf Handoff	VPLS over MPLS border leaf handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with a VPLS over MPLS network through a border leaf switch.

Release	Feature	Feature Information
Cisco IOS XE Amsterdam 17.1.1	Spine Functionality	A spine switch acts as the connecting node between all the leaf switches in a BGP EVPN VXLAN fabric, forwards the traffic between the leaf switches and provides redundancy to the network.
	Border Spine Functionality	A border spine switch in a BGP EVPN VXLAN fabric enables external connectivity with other Layer 2 and Layer 3 networks by acting as the connecting node between the two networks.
	Layer 3 Tenant Routed Multicast for IPv4 Traffic	<p>Layer 3 tenant routed multicast (TRM) for IPv4 traffic enables multicast forwarding for IPv4 traffic in a BGP EVPN VXLAN fabric. It provides multi-tenancy-aware multicast forwarding between senders and receivers within the same subnet or different subnets, locally or across VTEPs.</p> <p>You can configure TRM with PIM sparse mode (PIM-SM) using anycast RP mode and TRM with PIM source specific mode (PIM-SSM).</p>
	VRF-Lite Border Spine Handoff	VRF-Lite border spine handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with a VRF-Lite network through a border spine switch.
	MPLS Layer 3 VPN Border Spine Handoff	MPLS Layer 3 VPN border spine handoff in a BGP EVPN VXLAN fabric allows Layer 3 external connectivity with an MPLS Layer 3 VPN network through a border spine switch.
	IEEE 802.1Q Border Spine Handoff	IEEE 802.1Q border spine handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an IEEE 802.1Q network through a border spine switch.
	Access Network Border Spine Handoff	Access border spine handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with an Access network through a border spine switch.
	VPLS over MPLS Border Spine Handoff	VPLS over MPLS border spine handoff in a BGP EVPN VXLAN fabric allows Layer 2 external connectivity with a VPLS over MPLS network through a border spine switch.
	VXLAN-Aware Flexible Netflow	VXLAN-aware Flexible Netflow captures the VXLAN flow information for both bridged and routed traffic.
BGP EVPN VXLAN MIB support	Support was introduced for the MIB.	

Release	Feature	Feature Information
Cisco IOS XE Amsterdam 17.2.1	ARP and IPv6 Neighbor Discovery Flooding Suppression	Flooding suppression avoids the flooding of ARP and IPv6 neighbor discovery packets over the VXLAN network to the local and remote host or access devices.
	EVPN VXLAN Centralized Default Gateway	EVPN VXLAN Centralized Default Gateway allows a single VTEP in the EVPN VXLAN network to act as the Layer 3 gateway for all the Layer 2 VNIs in the network.
Cisco IOS XE Amsterdam 17.3.1	Layer 3 Tenant Routed Multicast for IPv6 Traffic	Layer 3 TRM for IPv6 traffic enables multicast forwarding for IPv6 traffic in a BGP EVPN VXLAN fabric. It provides multi-tenancy-aware multicast forwarding between senders and receivers within the same subnet or different subnets, locally or across VTEPs. Support was introduced to configure TRM with PIM-SSM and TRM with PIM-SM.
	Enhanced RP Functionality for Layer 3 TRM for IPv4 and IPv6 traffic	Enhancement was introduced for the configuration of an RP in the overlay network for TRM with PIM-SM. The enhancement allows you to configure an RP on a single or multiple VTEPs inside the BGP EVPN VXLAN fabric or on a device outside the fabric.
	Interworking of Layer 3 TRM with MVPN Networks for IPv4 Traffic	Interworking of Layer 3 TRM with MVPN Networks allows you to forward IPv4 Layer 3 multicast traffic between sources and receivers of an EVPN VXLAN network and an MVPN network.
	Broadcast, Unknown Unicast, and Multicast Traffic Rate Limiting	BUM Traffic Rate Limiting allows you to use a policer and set the flood rate limit of the BUM traffic in the network to a predefined value.
Cisco IOS XE Bengaluru 17.6.1	Private VLANs (PVLANS) in a BGP EVPN VXLAN Fabric	BGP EVPN VXLAN allows you to extend a PVLAN across the VTEPs in an EVPN fabric overlay. The extension uses isolated and community VLANs to provide features like subnet management and sub segmentation of a broadcast domain in the fabric.
Cisco IOS XE Bengaluru 17.6.1	EVPN VXLAN Multi-homing in Single-active Redundancy Mode	Multi-homing provides redundancy in the connection between a customer edge (CE) device and a VTEP by connecting the customer network with multiple VTEPs in an EVPN VXLAN network. In single-active redundancy mode, only one VTEP, among a group of VTEPs that are attached to the particular ethernet segment, is allowed to forward traffic to and from that ethernet segment. Multi-homing in single-active redundancy mode was introduced only in the form of dual-homing, allowing a CE device to be connected to two VTEPs.

Release	Feature	Feature Information
Cisco IOS XE Bengaluru 17.6.1	Route Leaking between EVPN and Global Routing Table	The route leaking feature provides the import and export of routes between an EVPN-enabled VRF and the Global Routing Table (GRT). By enabling the import and export of the EVPN learned routes into the GRT, the border node provides access to services such as DNS or DHCP to the hosts in the fabric.
Cisco IOS XE Bengaluru 17.6.2	Data MDT Support for L3 TRM	Data MDTs are purpose built underlay MDTs to provide optimized forwarding in the MVPN and EVPN core.
Cisco IOS XE Cupertino 17.8.1	Optimized Layer 2 Overlay Multicast for IPv4 and IPv6 traffic	Optimized Layer 2 Overlay Multicast forwards multicast traffic within the Layer 2 Virtual Network Instance (L2VNI). IPv4 and IPv6 Layer 2 overlay multicast traffic is optimized both at the access level and within the VXLAN fabric.
Cisco IOS XE Cupertino 17.9.1	BGP EVPN VXLAN over IPsec	An IPsec-based underlay network securely transports the VXLAN-encapsulated packets between the source and destination VTEPs. Securing the BGP EVPN VXLAN data traffic using IPsec tunnel encrypts the data and maintains data integrity. Support for this feature was introduced only on the Cisco Catalyst 9300X Series Switches.
Cisco IOS XE Dublin 17.10.1	BGP EVPN VXLAN with IPv6 in the Underlay (VXLANv6)	Introduced support for IPv6 addressing in the underlay of a BGP EVPN VXLAN fabric. In a new deployment, you can build your BGP EVPN VXLAN fabric with IPv6 underlay. For an existing BGP EVPN VXLAN fabric with IPv4 underlay, you can seamlessly migrate to an IPv6 or dual stack underlay.

Release	Feature	Feature Information
Cisco IOS XE Dublin 17.11.1	EVPN Microsegmentation	BGP EVPN VXLAN integrates Cisco TrustSec to provide microsegmentation and end-to-end access control with the propagation of the security group tag (SGT). Using security group-based access control lists (SGACLs), you can control the operations that a user can perform, based on the security group assignments and destination resources in a VXLAN campus fabric.
	EVPN VXLAN Multi-Homing	BGP EVPN is enhanced to restrict the ethernet segment operations to the EVPN-controlled VLANs on the trunk port. This allows traditional Layer 2 domains to co-exist with the Layer 2 VNI-enabled VLANs at access layer. It also allows selective VLAN migration to overlay VXLAN segmentation.
	Dynamic BGP Peering for EVPN	Support for BGP dynamic neighbor sessions is extended to the L2VPN EVPN address family.
	EVPN Route Map Support	The Leaf, Spine and Border nodes of a BGP EVPN fabric now support route map for the L2VPN address family. With 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.
Cisco IOS XE Dublin 17.11.99sw	TRM over BGP EVPN VXLANv6	TRM over BGP EVPN VXLANv6 follows the industry standard IETF RFCs 6513 and 6514 to support rich multicast applications over the BGP EVPN VXLAN fabric network. The Cisco Catalyst 9000 series switches provide flexibility to transport IPv4 multicast applications in overlay networks while the underlay network can be built on single-stack IPv6 only.
Cisco IOS XE Dublin 17.12.1	BGP EVPN VRF Auto RD and Auto RT	BGP EVPN Layer 3 overlay VRF configuration is simplified with the introduction of new CLIs to auto generate the route distinguisher (RD) and route target (RT) for a VRF.
	Dynamic ARP inspection (DAI) and DHCP Rogue Server Protection in BGP EVPN VXLAN fabric	BGP EVPN VXLAN fabric now supports ARP inspection and DHCP Rogue Server Protection. To configure these features, enable ARP inspection and DHCP Snooping on the VTEPs in the corresponding VLANs of the EVPN VXLAN fabric.

Release	Feature	Feature Information
Cisco IOS XE Dublin 17.12.2	TRM over BGP EVPN VXLANv6	<p>TRM over BGP EVPN VXLANv6 follows the industry standard IETF RFCs 6513 and 6514 to support rich multicast applications over the BGP EVPN VXLAN fabric network.</p> <p>The Cisco Catalyst 9000 series switches provide flexibility to transport IPv4 multicast applications in overlay networks while the underlay network can be built on single-stack IPv6 only.</p>

Use Cisco Feature Navigator to find information about platform and software image support. To access Cisco Feature Navigator, go to [Cisco Feature Navigator](#).

