Multihoming in a BGP EVPN VXLAN Fabric

Cisco IOS XE 26
Updated April 10, 2026

Full Cisco Trademarks with Software License

Preface

Preface

This preface describes the conventions of this document and information on how to obtain other documentation. It also provides information on what's new in Cisco product documentation.

Document Conventions

This document uses the following conventions:

<table>
<tr><th>Convention</th><th>Description</th></tr>
<tr><td>^ or Ctrl</td><td>Both the ^ symbol and Ctrl represent the Control (Ctrl) key on a keyboard. For example, the key combination **^D** or **Ctrl-D** means that you hold down the Control key while you press the D key. (Keys are indicated in capital letters but are not case sensitive.)</td></tr>
<tr><td>**bold** font</td><td>Commands and keywords and user-entered text appear in **bold** font.</td></tr>
<tr><td>*Italic* font</td><td>Document titles, new or emphasized terms, and arguments for which you supply values are in *italic* font.</td></tr>
<tr><td>`Courier` font</td><td>Terminal sessions and information the system displays appear in `courier` font.</td></tr>
<tr><td>**`Bold Courier`** font</td><td>**`Bold Courier`** font indicates text that the user must enter.</td></tr>
<tr><td>[x]</td><td>Elements in square brackets are optional.</td></tr>
<tr><td>...</td><td>An ellipsis (three consecutive nonbolded periods without spaces) after a syntax element indicates that the element can be repeated.</td></tr>
<tr><td>|</td><td>A vertical line, called a pipe, indicates a choice within a set of keywords or arguments.</td></tr>
<tr><td>[x | y]</td><td>Optional alternative keywords are grouped in brackets and separated by vertical bars.</td></tr>
<tr><td>{x | y}</td><td>Required alternative keywords are grouped in braces and separated by vertical bars.</td></tr>
<tr><td>[x {y | z}]</td><td>Nested set of square brackets or braces indicate optional or required choices within optional or required elements. Braces and a vertical bar within square brackets indicate a required choice within an optional element.</td></tr>
<tr><td>string</td><td>A nonquoted set of characters. Do not use quotation marks around the string or the string will include the quotation marks.</td></tr>
<tr><td>< ></td><td>Nonprinting characters such as passwords are in angle brackets.</td></tr>
<tr><td>[ ]</td><td>Default responses to system prompts are in square brackets.</td></tr>
<tr><td>!, #</td><td>An exclamation point (!) or a pound sign (#) at the beginning of a line of code indicates a comment line.</td></tr>
</table>

Reader Alert Conventions

This document may use the following conventions for reader alerts:

Note:

Means *reader take note*. Notes contain helpful suggestions or references to material not covered in the manual.

Tip:

Means *the following information will help you solve a problem.*

Caution:

Means *reader be careful*. In this situation, you might do something that could result in equipment damage or loss of data.

Timesaver:

Means *the described action saves time.* You can save time by performing the action described in the paragraph.

Take note of the following general safety warnings:

Warning:

Statement 1071—Warning Definition

IMPORTANT SAFETY INSTRUCTIONS

Before you work on any equipment, be aware of the hazards involved with electrical circuitry and be familiar with standard practices for preventing accidents. Read the installation instructions before using, installing, or connecting the system to the power source. Use the statement number at the beginning of each warning statement to locate its translation in the translated safety warnings for this device.

SAVE THESE INSTRUCTIONS

Related Documentation

Note:

Before installing or upgrading the device, refer to the device release notes.

Obtaining Documentation and Submitting a Service Request

For information on obtaining documentation, submitting a service request, and gathering additional information, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

http://www.cisco.com/c/en/us/td/docs/general/whatsnew/whatsnew.html

Subscribe to the *What's New in Cisco Product Documentation* as a Really Simple Syndication (RSS) feed and set content to be delivered directly to your desktop using a reader application. The RSS feeds are a free service and Cisco currently supports RSS version 2.0.

Topics included

1 Single-Active Mode

Topics:

This chapter describes the single-active mode, provides a comparison of single-active and all-active modes, and how to configure single-active networks.

Restrictions for single-active mode

Outlines the restrictions that apply to single-active mode.

Consider these restrictions when implementing multihoming in BGP EVPN VXLAN networks:

- Multihoming 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.
- EVPN multihoming is not supported on Cisco Catalyst 9500X Series Switches and Cisco Catalyst 9600-Sup2.
- Combined mode with StackWise Virtual and EVPN Multihoming is not supported in Cisco Catalyst 9400, 9500, and 9600 Series Switches.
- EVPN multihoming for non-fabric and fabric networks with native VXLANv6 is not supported.
- EVPN multihoming fabric networks in centralized gateway role is not supported.
- EVPN multihoming ISSU from a previous release to Cisco IOS XE 17.18.2 is not supported on Cisco Catalyst 9400 and 9600 Series Switches with redundant supervisors.
- Anycast gateway (global or VRF-enabled) SVI with any First Hop Redundancy Protocol (FHRP) protocol is not supported.

Single-active mode

Describes the redundancy feature that connects hosts or Layer 2 switches to EVPN VXLAN networks through multiple connection paths.

Multihoming in a BGP EVPN VXLAN fabric is a redundancy feature that

- provides redundant connections between a host or Layer 2 switch and the EVPN VXLAN network
- connects hosts or Layer 2 switches to the EVPN VXLAN network through either single-homing or multihoming, and
- supports RFC 7432 and RFC 8365 for VXLAN encapsulation-based EVPN multihoming capabilities on Cisco Catalyst 9000 Series switches.

Single-homing

Describes single-homing configuration to connect a host or a Layer 2 switch to a single VTEP in the EVPN VXLAN network.

Single-homing is a network connectivity method that

- allows you to connect a host or Layer 2 switch to a single VTEP in the EVPN VXLAN network
- does not support redundancy in the connection between the host or access device and the VTEP, and
- results in connection loss when the active link breaks down, making single-homed topologies not always reliable and efficient.

Single-homed topology

The following figure shows a single-homed topology:

Figure 1: Single-Homed topology

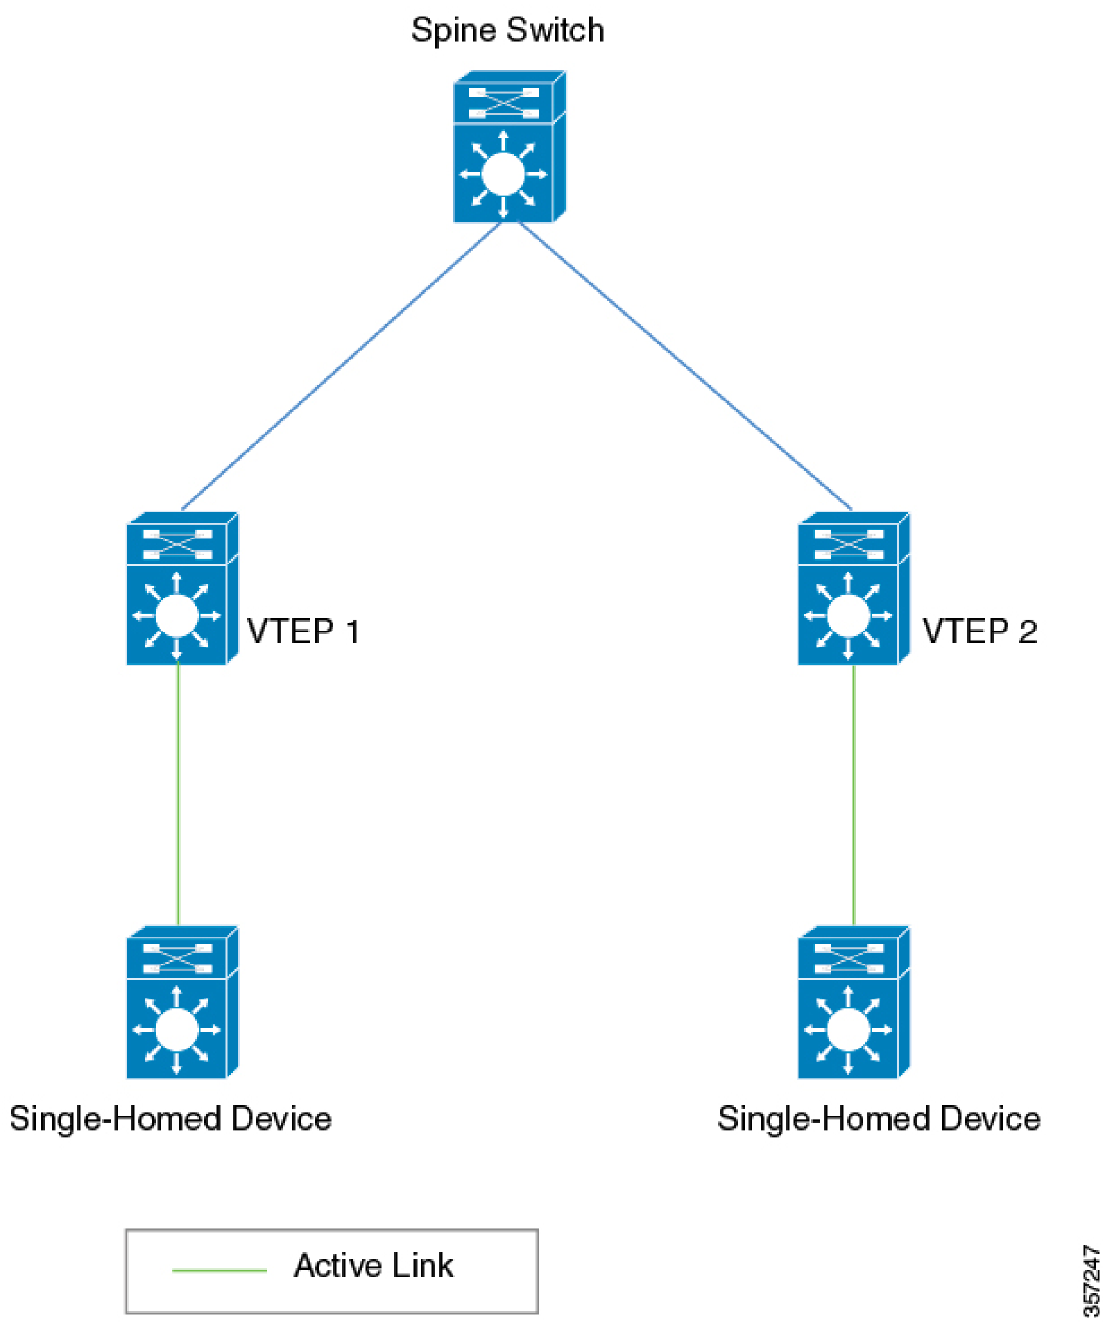

Multihoming

What is multihoming?

Multihoming allows you to connect a host or Layer 2 switch to more than one VTEP in a EVPN VXLAN network. This connection provides redundancy and allows network optimization. Redundancy in connection with VTEPs ensures that there is no traffic disruption when a network failure happens. Multihomed topologies are more reliant, secure, and efficient than single-homed topologies.

Multihoming operates in single-active and all-active redundancy modes. In both modes, the connected host or access device connects to the VTEPs through a set of links that together represent an Ethernet segment, identified by a globally unique Ethernet segment ID.

This Ethernet segment ID must also be part of the configuration on the VTEP 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.

ESI-MH in single-active redundancy mode

Describes single-active redundancy mode that results in a single-active access link between the VTEPs and the host (or Layer 2 switch) that passes through the Ethernet segment.

ESI-MH in single-active redundancy mode is an EVPN configuration that

- allows two VTEPs to act as an ESI-MH pair with only one VTEP forwarding traffic for a particular Ethernet-segment
- results in multiple VLANs on the physical link with a single link active for a given VLAN, and
- enables fabric-enabled VLANs and traditional Layer 2 VLANs to co-exist on the Ethernet-segment enabled trunk-downstream access interface from the VTEPs.

Configuration requirements

The single access link must be a physical Layer 2 interface. Physical interfaces can be part of a Layer 2 port channel for increased bandwidth, but member interfaces in the port channel must be connected to the same VTEP in single-active redundancy mode. From the downstream Layer 2 switch, there should be at least two independent port channel interfaces, with each of these interfaces terminating at a unique upstream VTEP.

Single-active redundancy mode is typically used when all the VLANs on a trunk port connected to an access switch cannot be migrated to an EVPN fabric.

Migration between single-homed and multihomed network topologies

Describes how BGP EVPN VXLAN enables network topology migration between different redundancy modes.

Migration between single-homed and multihomed network topologies is a BGP EVPN VXLAN capability that

- allows transition from a single-homed topology to a multihomed topology
- enables removal of redundancy from a multihomed topology to move back to a single-homed topology, and
- requires corresponding changes to the Ethernet segment configuration to prevent traffic loops and traffic blackholing.

Migration requirements

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 detailed migration procedures:

- For a sample illustration and detailed steps about how to migrate from a single-homed topology to a single-active dual-homed topology, see Migrate from a single-homed topology to a single-active dual-homed topology on page 15.
- For a sample illustration and detailed steps about how to migrate from a single-active dual-homed topology to a single-homed topology, see Migrate from a single-active dual-homed topology to a single-homed topology on page 17.

How to configure single-active mode

This section provides tasks to configure single-active mode in BGP EVPN VXLAN fabric networks.

Before you configure multihoming in a BGP EVPN VXLAN fabric, ensure that you configure EVPN VXLAN Layer 2 and Layer 3 overlay networks. For detailed steps, refer the How to Configure EVPN VXLAN Integrated Routing and Bridging section.

Configure dual-homing with single active redundancy in a BGP EVPN VXLAN fabric

Configure dual-homing with single active redundancy to provide network resilience in a BGP EVPN VXLAN fabric environment.

This task configures dual-homing with single-active redundancy in a BGP EVPN VXLAN fabric to provide network resilience and prevent traffic loops while ensuring only one path is active at a time.

Dual-homing with single-active redundancy is used in BGP EVPN VXLAN fabrics to provide redundant connectivity where only one of the redundant links is active at any given time, ensuring traffic follows a single path while maintaining backup connectivity.

Procedure

Perform the following set of procedures to configure dual-homing with single-active redundancy in a BGP EVPN VXLAN fabric.

Results

Dual-homing with single-active redundancy is configured in the BGP EVPN VXLAN fabric, providing redundant connectivity with loop prevention.

Configure redundancy on an Ethernet segment

Configure redundancy settings for an Ethernet segment in EVPN environments to ensure high availability and proper traffic forwarding.

Configure redundancy on an ethernet segment to provide high availability and ensure proper designated forwarder election in EVPN deployments.

Ethernet segment redundancy configuration is essential for EVPN multihoming scenarios where multiple provider edge devices connect to the same customer edge device or network segment.

Follow these steps to configure redundancy on an ethernet segment:

Procedure

1. **enable**

   ```
   Device> enable
   ```

   Enters privileged EXEC mode.

   Enter password, if prompted.

2. **configure terminal**

   ```
   Device# configure terminal
   ```

   Enters global configuration mode.

3. **l2vpn EVPN ethernet-segment** *ethernet-segment-id*

   ```
   Device(config)# l2vpn evpn ethernet-segment 1
   ```

Enters Layer 2 VPN EVPN ethernet segment configuration mode.

4. **identifier type** {**0** *ESI-value* | **3 system-MAC** *MAC-addrress*}

```
Device(config-evpn-es)# identifier type 0 0.0.0.0.0.0.0.0.1
```

Configures the ethernet segment identifier type (ESI) and value for the ethernet segment. The following ESI types are supported:

- Type 0: This type indicates an arbitrary 9-octet ESI value. The format is 00 + 9-octets of ESI value.
- Type 3: This type indicates a MAC-based ESI Value. The format is 03 + system-MAC (6 bytes) + value of MAC address (3 bytes).

5. **redundancy** *redundancy-type*

```
Device(config-evpn-es)# redundancy single-active
```

Configures the redundancy type for the ethernet segment.

6. **DF-election wait-time** *time-period*

```
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.

7. **end**

```
Device(config-evpn-es)# end
```

Exits Layer 2 VPN EVPN ethernet segment configuration mode and enters privileged EXEC mode.

Results

The ethernet segment redundancy configuration is complete. The device will participate in designated forwarder election based on the configured parameters and redundancy type.

Associate an Ethernet segment with an interface on a VTEP

Configure an Ethernet segment association with a VTEP interface for EVPN dual-homing functionality.

Associate an Ethernet segment with an interface on a VTEP to enable EVPN dual-homing functionality. This configuration allows multiple VTEPs to share connectivity to the same dual-homed device through a single Ethernet segment identifier.

This task is performed when configuring EVPN dual-homing where multiple VTEPs need to provide redundant connectivity to the same device. Each Ethernet segment requires a unique identifier that must be consistent across all VTEPs connecting to the same dual-homed device.

Procedure

1. **enable**

```
Device> enable
```

Enters privileged EXEC mode.

Enter password, if prompted.

2. **configure terminal**

```
Device# configure terminal
```

Enters global configuration mode.

3. **interface** *interface-ID*

```
Device(config)# interface GigabitEthernet1/0/10
```

Specifies the interface, and enters interface configuration mode.

4. **evpn ethernet-segment** *ethernet-segment-ID*

```
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).

5. **end**

```
Device(config-if)# end
```

Exits interface configuration mode and enters privileged EXEC mode.

Results

The Ethernet segment is now associated with the specified interface on the VTEP. The interface is configured for EVPN dual-homing with the assigned Ethernet segment ID.

Migrate from a single-homed topology to a single-active dual-homed topology

Configure the migration from a single-homed topology to a single-active dual-homed topology.

This task migrates a network from a single-homed configuration to a single-active dual-homed configuration to provide redundancy and improved network resilience.

When migrating between topologies, corresponding changes must be made to the Ethernet segment configuration. Failure to make corresponding changes to both the topology and Ethernet segment configuration results in traffic loops and traffic blackholing.

Figure 2: Migration from a Single-Homed network to a Single-Active Dual-Homed network

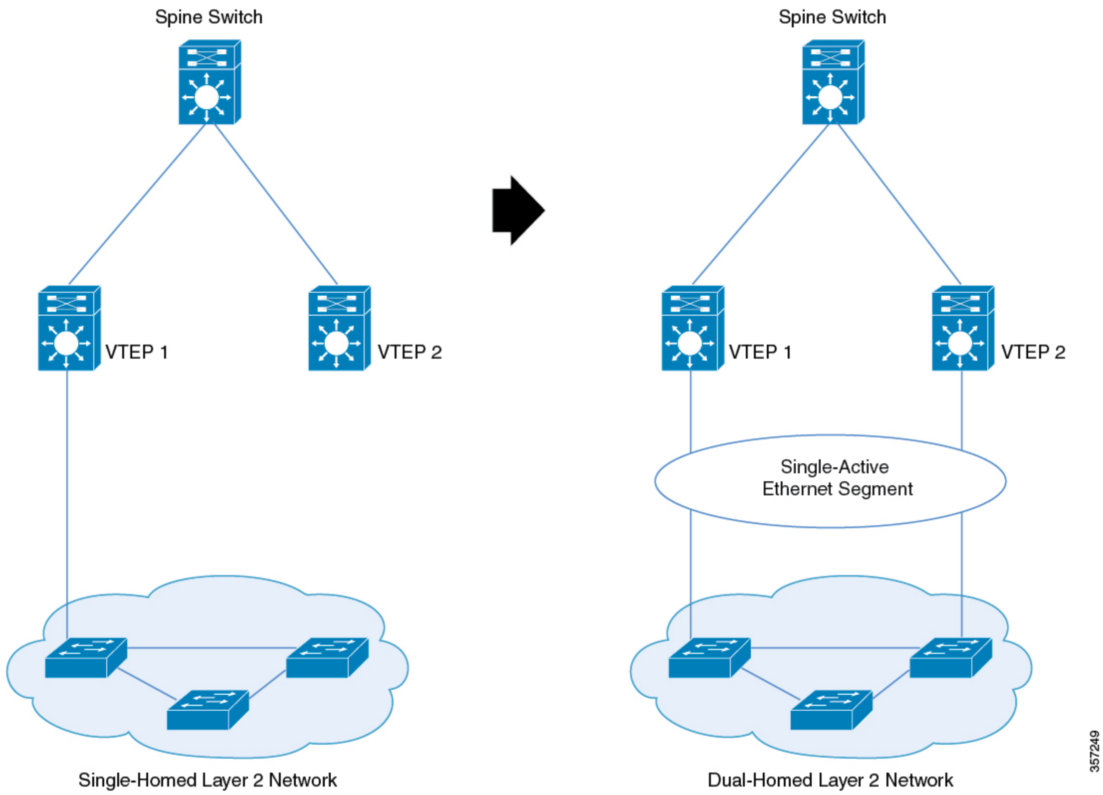

Before you migrate, we recommend that you do not configure the VTEP as the root bridge of the spanning tree, as the provision of Ethernet segment on the interface of a VTEP excludes it from spanning-tree. If the VTEP is the root bridge, its exclusion from the spanning-tree triggers an immediate spanning tree re-convergence.

Follow these steps to migrate from a single-homed topology to a single-active dual-homed topology:

Procedure

1. Ensure that you do not activate a link between VTEP 2 and a switch in the single-homed network.

   Activate the second link 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. 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.

Results

The single-homed network has now migrated to a dual-homed network.

Migrate from a single-active dual-homed topology to a single-homed topology

Configure a migration from a single-active dual-homed topology to a single-homed topology.

This task enables you to change your network topology from a single-active dual-homed configuration to a single-homed configuration while maintaining network connectivity and avoiding traffic disruptions.

The following figure illustrates the migration from a single-active dual-homed topology to a single-homed topology:

Figure 3: Migration from a Single-Active Dual-Homed network to a Single-Homed network

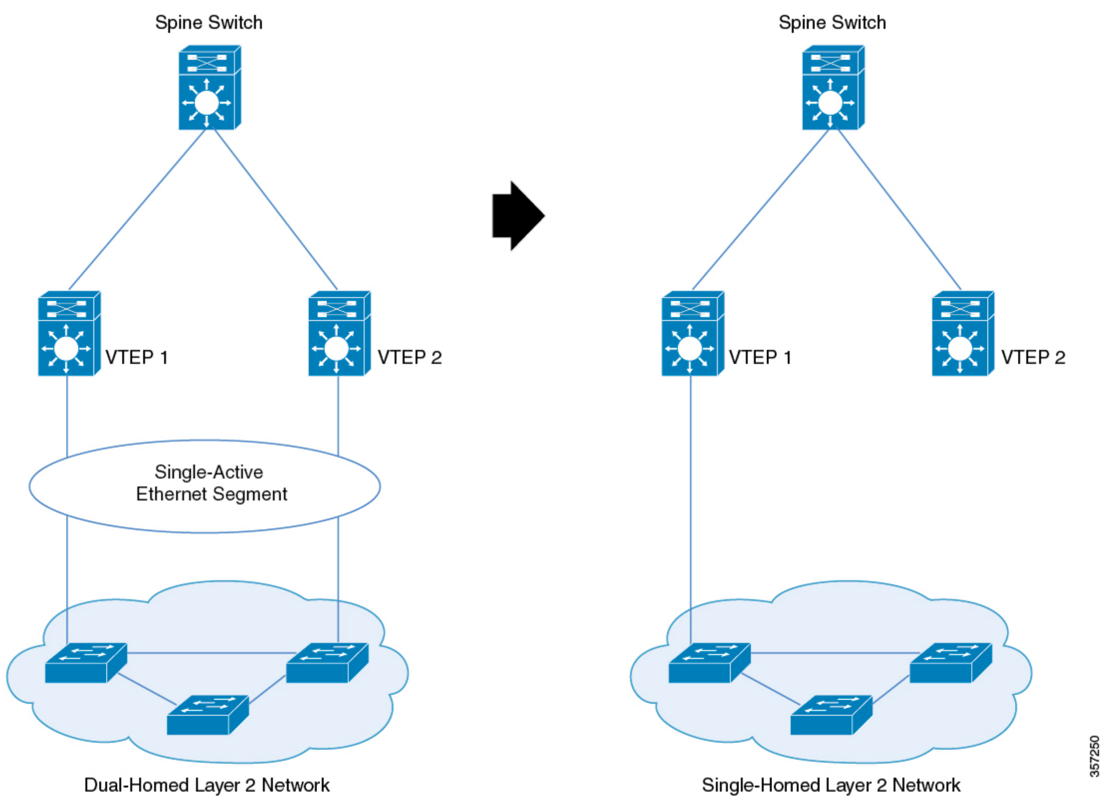

Note:

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.

Follow these steps to migrate from a single-active dual-homed network to a single-homed network:

Procedure

1. Configure portfast on the link you plan to activate.

   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.

Results

You have successfully migrated from a single-active dual-homed topology to a single-homed topology. The network now operates with a single connection path, and all traffic flows through the active link.

Configuration examples for multihoming in single-active redundancy mode

Describes configuration examples for implementing multihoming in single-active redundancy mode, providing practical setup guidance for network redundancy scenarios.

This section provides configuration examples for multihoming in single-active redundancy mode.

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

Sample configuration for 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 4: Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric

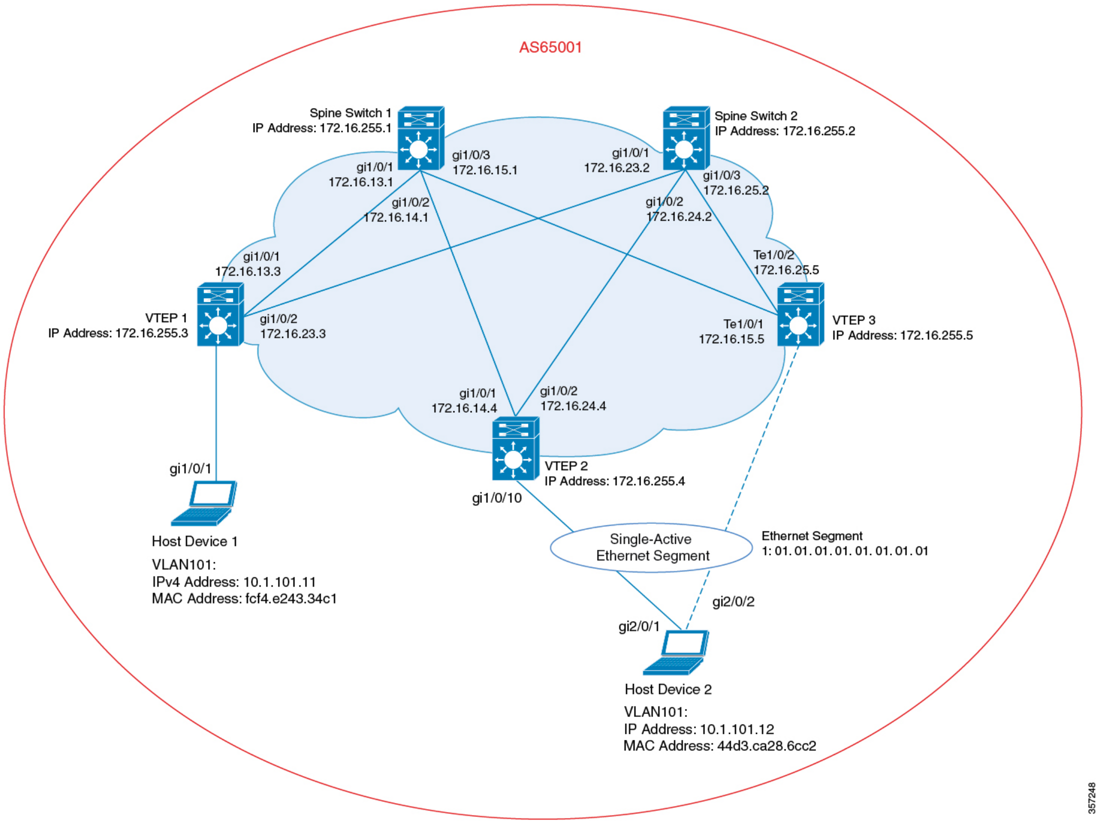

The topology shows an EVPN VXLAN network with two spine switches (Spine Switch 1 and Spine Switch 2) and three VTEPs (VTEP 1, VTEP 2, and VTEP 3). Host Device 1 is connected to VTEP 1. Host Device 2 is connected to VTEP 2 and VTEP 3 as a dual-homed single-active connection that passes through Ethernet 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 Verify Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric on page 26 section.

Table 1: Configuring Dual-Homing with Single-Active Redundancy using VTEP 2 and VTEP 3

VTEP 2	VTEP 3
Leaf-02# **show running-config** hostname Leaf-02 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn ethernet-segment 1 identifier type 0 01.01.01.01.01.01.01.01.01 redundancy single-active ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.4 255.255.255.255 ip ospf 1 area 0 !	LEaf-03# **show running-config** hostname Leaf-03 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! address-family ipv6 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! ip routing ! l2vpn evpn replication-type static router-id Loopback1 default-gateway advertise ! l2vpn evpn ethernet-segment 1 identifier type 0 01.01.01.01.01.01.01.01.01 redundancy single-active ! l2vpn evpn instance 101 vlan-based encapsulation vxlan replication-type ingress ! system mtu 9198 ! vlan configuration 101 member evpn-instance 101 vni 10101 vlan configuration 901 member vni 50901 ! interface Loopback0 ip address 172.16.255.5 255.255.255.255 ip ospf 1 area 0 !

VTEP 2

```
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
```

VTEP 3

```
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
```

VTEP 2	VTEP 3

`!`
`router ospf 1`
`router-id 172.16.255.4`
`!`
`router bgp 65001`
`bgp log-neighbor-changes`
`no bgp default ipv4-unicast`
`neighbor 172.16.255.1 remote-as 65001`
`neighbor 172.16.255.1 update-source Loopback0`
`neighbor 172.16.255.2 remote-as 65001`
`neighbor 172.16.255.2 update-source Loopback0`
`!`
`address-family ipv4`
`exit-address-family`
`!`
`address-family l2vpn evpn`
`neighbor 172.16.255.1 activate`
`neighbor 172.16.255.1 send-community both`
`neighbor 172.16.255.2 activate`
`neighbor 172.16.255.2 send-community both`
`exit-address-family`
`!`
`address-family ipv4 vrf green`
`advertise l2vpn evpn`
`redistribute connected`
`redistribute static`
`exit-address-family`
`!`
`end`
`!`
`Leaf-02#`
	`!`
	`router ospf 1`
	`router-id 172.16.255.5`
	`!`
	`router bgp 65001`
	`bgp log-neighbor-changes`
	`no bgp default ipv4-unicast`
	`neighbor 172.16.255.1 remote-as 65001`
	`neighbor 172.16.255.1 update-source Loopback0`
	`neighbor 172.16.255.2 remote-as 65001`
	`neighbor 172.16.255.2 update-source Loopback0`
	`!`
	`address-family ipv4`
	`exit-address-family`
	`!`
	`address-family l2vpn evpn`
	`neighbor 172.16.255.1 activate`
	`neighbor 172.16.255.1 send-community both`
	`neighbor 172.16.255.2 activate`
	`neighbor 172.16.255.2 send-community both`
	`exit-address-family`
	`!`
	`address-family ipv4 vrf green`
	`advertise l2vpn evpn`
	`redistribute connected`
	`redistribute static`
	`exit-address-family`
	`!`
	`end`
	`!`
	`Leaf-03#`

Table 2: Configuring Spine Switch 1, Spine Switch 2, and VTEP 1 to Configure Dual-Homing with Single-Active Redundancy

Spine Switch 1	Spine Switch 2	VTEP 1

Spine Switch 1	Spine Switch 2	VTEP 1
Spine-01# **show running-config** hostname Spine-01 ! ip routing ! system mtu 9198 ! interface Loopback0 ip address 172.16.255.1 255.255.255.255 ip ospf 1 area 0 ! interface 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 !	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 !	

Spine Switch 1	Spine Switch 2	VTEP 1
		Leaf-01# **show running-config** hostname Leaf-01 ! vrf definition green rd 1:1 ! address-family ipv4 route-target export 1:1 route-target import 1:1 route-target export 1:1 stitching route-target import 1:1 stitching exit-address-family ! 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

Spine Switch 1	Spine Switch 2	VTEP 1
		`ip ospf 1 area 0` `!`

Spine Switch 1

```
router bgp 65001
bgp router-id
172.16.255.1
bgp log-neighbor-changes
no bgp default ipv4-uni►
cast
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#
```

Spine Switch 2

```
router bgp 65001
bgp router-id
172.16.255.2
bgp log-neighbor-changes
no bgp default ipv4-uni►
cast
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#
```

VTEP 1

```
interface GigabitEther►
net0/0
vrf forwarding Mgmt-vrf
ip address 10.62.149.179
 255.255.255.0
negotiation auto
!
interface GigabitEther►
net1/0/1
no switchport
ip address 172.16.13.3
255.255.255.0
ip ospf network point-
to-point
ip ospf 1 area 0
!
interface GigabitEther►
net1/0/2
no switchport
ip address 172.16.23.3
255.255.255.0
ip ospf network point-
to-point
ip ospf 1 area 0
!
interface GigabitEther►
net1/0/10
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 Loop►
back1
host-reachability proto►
col bgp
member vni 10101
ingress-replication
member vni 50901 vrf
green
!
router ospf 1
router-id 172.16.255.3
!
```

Verify 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 27
- Outputs to Verify the Configuration on VTEP 2 on page 30
- Outputs to Verify the Configuration on VTEP 3 on page 33
- Outputs to Verify the Configuration on Spine Switch 1 on page 37
- Outputs to Verify the Configuration on Spine Switch 2 on page 40

Outputs to Verify the Configuration on VTEP 1

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

```
                  Leaf-01#
                  show nve peer
Interface   VNI      Type Peer-IP          RMAC/Num_RTs   eVNI      state flags
 UP time
nve1        50901    L3CP 172.16.254.5     7c21.0dbd.2748 50901       UP  A/M/4
 01:17:04
nve1        50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901       UP  A/M/4
 03:26:09
nve1        10101    L2CP 172.16.254.4     8              10101       UP   N/A
 03:52:15
nve1        10101    L2CP 172.16.254.5     10             10101       UP   N/A
 05:25:28
Leaf-01#
```

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

```
                  Leaf-01#
                  show l2vpn evpn evi 101 detail
EVPN instance:         101 (VLAN Based)
  RD:                  172.16.254.3:101 (auto)
  Import-RTs:          65001:101
  Export-RTs:          65001:101
  Per-EVI Label:       none
  State:               Established
  Replication Type:    Ingress
  Encapsulation:       vxlan
  IP Local Learn:      Enabled (global)
  Adv. Def. Gateway:   Enabled (global)
  Vlan:                101
    Ethernet-Tag:      0
    State:             Established
    Core If:           Vlan901
    Access If:         Vlan101
    NVE If:            nve1
    RMAC:              10b3.d56a.8fc8
    Core Vlan:         901
    L2 VNI:            10101
    L3 VNI:            50901
    VTEP IP:           172.16.254.3
    VRF:               green
```

```
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101
        Routes: 1 MAC, 1 MAC/IP
    Peers:
      172.16.254.4
        Routes: 4 MAC, 2 MAC/IP, 1 IMET, 1 EAD
      172.16.254.5
        Routes: 6 MAC, 2 MAC/IP, 1 IMET, 1 EAD
Leaf-01#
```

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

```
                Leaf-01#
                show bgp l2vpn evpn evi 101
BGP table version is 6958, local router ID is 172.16.255.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,

              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.3:101
 *>i  [1][172.16.254.3:101][00010101010101010101][0]/23
                      172.16.254.5             0    100      0 ?
 *mi                   172.16.254.4             0    100      0 ?
 *>   [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      ::                                 32768 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286C82][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.3:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      ::                                 32768 ?
 *>   [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      ::                                 32768 ?
 *>   [3][172.16.254.3:101][0][32][172.16.254.3]/17
                      ::                                 32768 ?
 *>i  [3][172.16.254.3:101][0][32][172.16.254.4]/17
                      172.16.254.4             0    100      0 ?
 *>i  [3][172.16.254.3:101][0][32][172.16.254.5]/17
                      172.16.254.5             0    100      0 ?
Leaf-01#
```

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

```
                Leaf-01#
                show l2route evpn mac
  EVI       ETag  Prod    Mac Address                  Next Hop(s) Seq Number
```

```
----- ---------- ----- -------------- ---------------------------- ----------

  101          0 L2VPN 10b3.d56a.8fc1                      Vl101:0          0

  101          0   BGP 44d3.ca28.6c82        V:10101 172.16.254.5          0

  101          0   BGP 44d3.ca28.6cc2        V:10101 172.16.254.5          0

  101          0   BGP 7c21.0dbd.2741        V:10101 172.16.254.5          0

  101          0   BGP 7c21.0dbd.9541        V:10101 172.16.254.4          0

  101          0 L2VPN f4cf.e243.34c1                 Gi1/0/10:101          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 Verify Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric on page 26 .

Outputs to Verify the Configuration on VTEP 2

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

```
                Leaf-02#
                show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI     state flags
 UP time
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/4
 03:24:45
nve1       50901    L3CP 172.16.254.5     7c21.0dbd.2748 50901      UP  A/M/4
 01:15:39
nve1       10101    L2CP 172.16.254.3     5              10101      UP   N/A
 03:24:45
nve1       10101    L2CP 172.16.254.5     6              10101      UP   N/A
 03:24:45
Leaf-02#
```

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

```
                Leaf-02#
                show l2vpn evpn ethernet-segment detail
EVPN Ethernet Segment ID: 0001.0101.0101.0101.0101
  Interface:              Gi1/0/10
  Redundancy mode:        single-active
  DF election wait time:  3 seconds
  Split Horizon label:    0
  State:                  Ready
  Encapsulation:          vxlan
  Ordinal:                0
  RD:                     172.16.254.4:7
    Export-RTs:           65001:101
  Forwarder List:         172.16.254.4 172.16.254.5
Leaf-02#
```

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

```
                Leaf-02#
                show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.4:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Vlan:              101
    Ethernet-Tag:    0
    State:           Established
    Core If:         Vlan901
    Access If:       Vlan101
```

```
    NVE If:          nve1
    RMAC:            7c21.0dbd.9548
    Core Vlan:       901
    L2 VNI:          10101
    L3 VNI:          50901
    VTEP IP:         172.16.254.4
    VRF:             green
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101 (DF state: blocked)
        Routes: 0 MAC, 0 MAC/IP
    Peers:
      172.16.254.3
        Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
      172.16.254.5
        Routes: 3 MAC, 1 MAC/IP, 1 IMET, 1 EAD
Leaf-02#
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 4 on VTEP 2:

```
                Leaf-02#
                show bgp l2vpn evpn route-type 4
BGP routing table entry for
[4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23, version 601
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
     1
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.4)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local,
best
      Local vtep: 172.16.254.4
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:41:40 UTC
BGP routing table entry for
[4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23, version 658
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 6
  Local
    172.16.254.5 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.5, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Jan 26 2021 19:43:19 UTC
  Refresh Epoch 6
  Local
    172.16.254.5 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.5, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:43:19 UTC
Leaf-02#
```

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

```
                Leaf-02#
                show bgp l2vpn evpn evi 101
BGP table version is 845, local router ID is 172.16.255.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,

              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.4:101
 *mi  [1][172.16.254.4:101][00010101010101010101][0]/23
                      172.16.254.5             0    100      0 ?
 *>                    ::                                 32768 ?
 *>i  [2][172.16.254.4:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286C82][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                      172.16.254.5             0    100      0 ?
 *>   [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      ::                                 32768 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.4:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      172.16.254.3             0    100      0 ?
 *>i  [3][172.16.254.4:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>   [3][172.16.254.4:101][0][32][172.16.254.4]/17
                      ::                                 32768 ?
 *>i  [3][172.16.254.4:101][0][32][172.16.254.5]/17
                      172.16.254.5             0    100      0 ?
Leaf-02#
```

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

```
                Leaf-02#
                show l2route evpn mac
  EVI       ETag  Prod    Mac Address                    Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------

  101          0   BGP 10b3.d56a.8fc1         V:10101 172.16.254.3          0

  101          0   BGP 44d3.ca28.6c82         V:10101 172.16.254.5          0

  101          0   BGP 44d3.ca28.6cc2         V:10101 172.16.254.5          0

  101          0   BGP 7c21.0dbd.2741         V:10101 172.16.254.5          0

  101          0 L2VPN 7c21.0dbd.9541                     Vl101:0          0

  101          0   BGP f4cf.e243.34c1         V:10101 172.16.254.3          0
```

```
Leaf-02#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* command on VTEP 2:

```
                Leaf-02#
                show l2route evpn mac esi 0001.0101.0101.0101.0101
  EVI       ETag  Prod    Mac Address                  Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------

  101          0   BGP 44d3.ca28.6c82         V:10101 172.16.254.5          0

  101          0   BGP 44d3.ca28.6cc2         V:10101 172.16.254.5          0
Leaf-02#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* **detail** command on VTEP 2:

```
                Leaf-02#
                show l2route evpn mac esi 0001.0101.0101.0101.0101 detail
EVPN Instance:             101
Ethernet Tag:              0
Producer Name:             BGP
MAC Address:               44d3.ca28.6c82
Num of MAC IP Route(s):    0
Sequence Number:           0
ESI:                       0001.0101.0101.0101.0101
Flags:                     B()
Next Hop(s):               V:10101 172.16.254.5
Resolved Next Hops:        V:10101 172.16.254.5
Resolved Redundancy Mode: Single-Active
EVPN Instance:             101
Ethernet Tag:              0
Producer Name:             BGP
MAC Address:               44d3.ca28.6cc2
Num of MAC IP Route(s):    0
Sequence Number:           0
ESI:                       0001.0101.0101.0101.0101
Flags:                     B()
Next Hop(s):               V:10101 172.16.254.5
Resolved Next Hops:        V:10101 172.16.254.5
Resolved Redundancy Mode: Single-Active
Leaf-02#
```

Return to Verify Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric on page 26 .

Outputs to Verify the Configuration on VTEP 3

The following example shows the output for the **show nve peer** command on VTEP 3:

```
                Leaf-03#
                show nve peer
Interface  VNI      Type Peer-IP          RMAC/Num_RTs   eVNI      state flags
 UP time
```

```
nve1       50901    L3CP 172.16.254.3     10b3.d56a.8fc8 50901      UP  A/M/4
 04:23:46
nve1       50901    L3CP 172.16.254.4     7c21.0dbd.9548 50901      UP  A/M/4
 03:24:57
nve1       10101    L2CP 172.16.254.3     5              10101      UP   N/A
 04:23:46
nve1       10101    L2CP 172.16.254.4     4              10101      UP   N/A
 03:24:57
Leaf-03#
```

The following example shows the output for the **show l2vpn evpn ethernet-segment detail** command on VTEP 3:

```
                Leaf-03#
                show l2vpn evpn ethernet-segment detail
EVPN Ethernet Segment ID: 0001.0101.0101.0101.0101
  Interface:              Gi1/0/10
  Redundancy mode:        single-active
  DF election wait time:  3 seconds
  Split Horizon label:    0
  State:                  Ready
  Encapsulation:          vxlan
  Ordinal:                1
  RD:                     172.16.254.5:9
    Export-RTs:           65001:101
  Forwarder List:         172.16.254.4 172.16.254.5
Leaf-03#
```

The following example shows the output for the **show l2vpn evpn evi** *evpn-instance* **detail** command on VTEP 3:

```
                Leaf-03#
                show l2vpn evpn evi 101 detail
EVPN instance:       101 (VLAN Based)
  RD:                172.16.254.5:101 (auto)
  Import-RTs:        65001:101
  Export-RTs:        65001:101
  Per-EVI Label:     none
  State:             Established
  Replication Type:  Ingress
  Encapsulation:     vxlan
  IP Local Learn:    Enabled (global)
  Adv. Def. Gateway: Enabled (global)
  Vlan:              101
    Ethernet-Tag:    0
    State:           Established
    Core If:         Vlan901
    Access If:       Vlan101
    NVE If:          nve1
    RMAC:            7c21.0dbd.2748
    Core Vlan:       901
    L2 VNI:          10101
    L3 VNI:          50901
    VTEP IP:         172.16.254.5
    VRF:             green
    IPv4 IRB:        Enabled
    IPv6 IRB:        Disabled
    Pseudoports:
      GigabitEthernet1/0/10 service instance 101 (DF state: forwarding)
```

```
         Routes: 2 MAC, 0 MAC/IP
     Peers:
       172.16.254.3
         Routes: 2 MAC, 2 MAC/IP, 1 IMET, 0 EAD
       172.16.254.4
         Routes: 1 MAC, 1 MAC/IP, 1 IMET, 1 EAD
Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn route-type** command for route type 4 on VTEP 3:

```
                Leaf-03#
                show bgp l2vpn evpn route-type 4
BGP routing table entry for
[4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23, version 337
Paths: (2 available, best #2, table EVPN-BGP-Table)
  Not advertised to any peer
  Refresh Epoch 5
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.2 (172.16.255.2)
      Origin incomplete, metric 0, localpref 100, valid, internal
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.4, Cluster list: 172.16.255.2
      rx pathid: 0, tx pathid: 0
      Updated on Jan 26 2021 19:38:35 UTC
  Refresh Epoch 5
  Local
    172.16.254.4 (metric 3) (via default) from 172.16.255.1 (172.16.255.1)
      Origin incomplete, metric 0, localpref 100, valid, internal, best
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      Originator: 172.16.255.4, Cluster list: 172.16.255.1
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:38:35 UTC
BGP routing table entry for
[4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23, version 1269
Paths: (1 available, best #1, table EVPN-BGP-Table)
  Advertised to update-groups:
     2
  Refresh Epoch 1
  Local
    :: (via default) from 0.0.0.0 (172.16.255.5)
      Origin incomplete, localpref 100, weight 32768, valid, sourced, local,
best
      Local vtep: 172.16.254.5
      Extended Community: ENCAP:8 EVPN ES-IMPORT:0x101:0x101:0x101
      rx pathid: 0, tx pathid: 0x0
      Updated on Jan 26 2021 19:40:14 UTC
Leaf-03#
```

The following example shows the output for the **show bgp l2vpn evpn evi** *evpn-instance* command on VTEP 3:

```
                Leaf-03#
                show bgp l2vpn evpn evi 101
BGP table version is 1284, local router ID is 172.16.255.5
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
```

```
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.5:101
 *>   [1][172.16.254.5:101][00010101010101010101][0]/23
                      ::                                 32768 ?
 *mi                   172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.5:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>   [2][172.16.254.5:101][0][48][44D3CA286C82][0][*]/20
                      ::                                 32768 ?
 *>   [2][172.16.254.5:101][0][48][44D3CA286CC2][0][*]/20
                      ::                                 32768 ?
 *>   [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                      ::                                 32768 ?
 *>i  [2][172.16.254.5:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i  [2][172.16.254.5:101][0][48][F4CFE24334C1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i  [2][172.16.254.5:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      172.16.254.3             0    100      0 ?
 *>i  [3][172.16.254.5:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>i  [3][172.16.254.5:101][0][32][172.16.254.4]/17
                      172.16.254.4             0    100      0 ?
 *>   [3][172.16.254.5:101][0][32][172.16.254.5]/17
                      ::                                 32768 ?
Leaf-03#
```

The following example shows the output for the **show l2route evpn mac** command on VTEP 3:

```
                Leaf-03#
                show l2route evpn mac
  EVI       ETag  Prod    Mac Address                  Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------

  101          0   BGP 10b3.d56a.8fc1        V:10101 172.16.254.3          0

  101          0 L2VPN 44d3.ca28.6c82                  Gi1/0/10:101          0

  101          0 L2VPN 44d3.ca28.6cc2                  Gi1/0/10:101          0

  101          0 L2VPN 7c21.0dbd.2741                       Vl101:0          0

  101          0   BGP 7c21.0dbd.9541        V:10101 172.16.254.4          0

  101          0   BGP f4cf.e243.34c1        V:10101 172.16.254.3          0
Leaf-03#
```

The following example shows the output for the **show l2route evpn mac esi** *ethernet-segment-id* command on VTEP 3:

```
                Leaf-03#
                show l2route evpn mac esi 0001.0101.0101.0101.0101
```

```
  EVI       ETag  Prod    Mac Address                  Next Hop(s) Seq Number
----- ---------- ----- -------------- ---------------------------- ----------

  101          0 L2VPN 44d3.ca28.6c82                 Gi1/0/10:101          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
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 Verify Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric on page 26 .

Outputs to Verify the Configuration on Spine Switch 1

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

```
                Spine-01#
                show bgp l2vpn evpn summary
BGP router identifier 172.16.255.1, local AS number 65001
BGP table version is 5443, main routing table version 5443
17 network entries using 5848 bytes of memory
34 path entries using 7072 bytes of memory
13/11 BGP path/bestpath attribute entries using 3744 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
10 BGP extended community entries using 480 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 17264 total bytes of memory
BGP activity 101/84 prefixes, 2825/2791 paths, scan interval 60 secs
25 networks peaked at 14:54:41 Jan 26 2021 UTC (05:39:56.356 ago)
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
172.16.255.2    4        65001    5664    5668     5443    0    0 05:40:29
    15
```

```
172.16.255.3    4         65001     378    5690     5443    0    0 05:35:23
     5
172.16.255.4    4         65001     440    1633     5443    0    0 03:36:33
     6
172.16.255.5    4         65001     594    5296     5443    0    0 04:34:27
     8
Spine-01#
```

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

```
                Spine-01#
                show bgp l2vpn evpn
BGP table version is 5443, local router ID is 172.16.255.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,

              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.4:7
 *>i  [1][172.16.254.4:7][00010101010101010101][4294967295]/23
                      172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [1][172.16.254.4:101][00010101010101010101][0]/23
                      172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:9
 *>i  [1][172.16.254.5:9][00010101010101010101][4294967295]/23
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 *>i  [1][172.16.254.5:101][00010101010101010101][0]/23
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [2][172.16.254.5:101][0][48][44D3CA286C82][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
```

```
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [3][172.16.254.5:101][0][32][172.16.254.5]/17
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:257
 * i  [4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:257
 * i  [4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                      172.16.254.5             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i                   172.16.254.3             0    100      0 ?
Spine-01#
```

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

```
                Spine-01#
                show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route

       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 17 subnets, 2 masks
C        172.16.13.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.13.1/32 is directly connected, GigabitEthernet1/0/1
C        172.16.14.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.14.1/32 is directly connected, GigabitEthernet1/0/2
C        172.16.15.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.15.1/32 is directly connected, GigabitEthernet1/0/3
O        172.16.23.0/24
           [110/2] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.24.0/24
           [110/2] via 172.16.14.4, 03:37:00, GigabitEthernet1/0/2
O        172.16.25.0/24
           [110/2] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
O        172.16.254.3/32
           [110/2] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
```

```
O        172.16.254.4/32
           [110/2] via 172.16.14.4, 03:36:50, GigabitEthernet1/0/2
O        172.16.254.5/32
           [110/2] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
C        172.16.255.1/32 is directly connected, Loopback0
O        172.16.255.2/32
           [110/3] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
           [110/3] via 172.16.14.4, 03:37:00, GigabitEthernet1/0/2
           [110/3] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.255.3/32
           [110/2] via 172.16.13.3, 05:35:46, GigabitEthernet1/0/1
O        172.16.255.4/32
           [110/2] via 172.16.14.4, 03:36:56, GigabitEthernet1/0/2
O        172.16.255.5/32
           [110/2] via 172.16.15.5, 03:38:33, GigabitEthernet1/0/3
Spine-01#
```

Return to Verify Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric on page 26 .

Outputs to Verify the Configuration on Spine Switch 2

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

```
                Spine-02#
                show bgp l2vpn evpn summary
BGP router identifier 172.16.255.2, local AS number 65001
BGP table version is 5499, main routing table version 5499
17 network entries using 5848 bytes of memory
34 path entries using 7072 bytes of memory
13/11 BGP path/bestpath attribute entries using 3744 bytes of memory
3 BGP rrinfo entries using 120 bytes of memory
10 BGP extended community entries using 480 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 17264 total bytes of memory
BGP activity 101/84 prefixes, 2823/2789 paths, scan interval 60 secs
25 networks peaked at 14:56:03 Jan 26 2021 UTC (05:40:54.652 ago)
Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down
State/PfxRcd
172.16.255.1    4        65001    5669    5665     5499    0    0 05:41:28
    15
172.16.255.3    4        65001     381    5691     5499    0    0 05:36:22
     5
172.16.255.4    4        65001     440    1632     5499    0    0 03:37:31
     6
172.16.255.5    4        65001     594    5291     5499    0    0 04:35:26
     8
Spine-02#
```

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

```
                Spine-02#
                show bgp l2vpn evpn
BGP table version is 5499, local router ID is 172.16.255.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,

             r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
```

```
              x best-external, a additional-path, c RIB-compressed,
              t secondary path, L long-lived-stale,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found
     Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 172.16.254.4:7
 *>i  [1][172.16.254.4:7][00010101010101010101][4294967295]/23
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 *>i  [1][172.16.254.4:101][00010101010101010101][0]/23
                      172.16.254.4             0    100      0 ?
 * i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:9
 *>i  [1][172.16.254.5:9][00010101010101010101][4294967295]/23
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 *>i  [1][172.16.254.5:101][00010101010101010101][0]/23
                      172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [2][172.16.254.3:101][0][48][10B3D56A8FC1][32][10.1.101.1]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][0][*]/20
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
 * i  [2][172.16.254.3:101][0][48][F4CFE24334C1][32][10.1.101.11]/24
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [2][172.16.254.4:101][0][48][7C210DBD9541][32][10.1.101.1]/24
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [2][172.16.254.5:101][0][48][44D3CA286C82][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][44D3CA286CC2][0][*]/20
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
 * i  [2][172.16.254.5:101][0][48][7C210DBD2741][32][10.1.101.1]/24
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.254.3:101
 * i  [3][172.16.254.3:101][0][32][172.16.254.3]/17
                      172.16.254.3             0    100      0 ?
 *>i                   172.16.254.3             0    100      0 ?
Route Distinguisher: 172.16.254.4:101
 * i  [3][172.16.254.4:101][0][32][172.16.254.4]/17
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.254.5:101
 * i  [3][172.16.254.5:101][0][32][172.16.254.5]/17
                      172.16.254.5             0    100      0 ?
 *>i                   172.16.254.5             0    100      0 ?
Route Distinguisher: 172.16.255.4:257
 * i  [4][172.16.255.4:257][00010101010101010101][32][172.16.254.4]/23
                      172.16.254.4             0    100      0 ?
 *>i                   172.16.254.4             0    100      0 ?
Route Distinguisher: 172.16.255.5:257
 * i  [4][172.16.255.5:257][00010101010101010101][32][172.16.254.5]/23
```

```
                      172.16.254.5            0    100      0 ?
 *>i                   172.16.254.5            0    100      0 ?
Route Distinguisher: 1:1
 * i  [5][1:1][0][24][10.1.101.0]/17
                      172.16.254.5            0    100      0 ?
 * i                   172.16.254.4            0    100      0 ?
 *>i                   172.16.254.3            0    100      0 ?
 * i                   172.16.254.3            0    100      0 ?
Spine-02#
```

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

```
                Spine-02#
                show ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2, m - OMP
       n - NAT, Ni - NAT inside, No - NAT outside, Nd - NAT DIA
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route

       H - NHRP, G - NHRP registered, g - NHRP registration summary
       o - ODR, P - periodic downloaded static route, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
Gateway of last resort is not set
      172.16.0.0/16 is variably subnetted, 17 subnets, 2 masks
O        172.16.13.0/24
           [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
O        172.16.14.0/24
           [110/2] via 172.16.24.4, 03:37:38, GigabitEthernet1/0/2
O        172.16.15.0/24
           [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
C        172.16.23.0/24 is directly connected, GigabitEthernet1/0/1
L        172.16.23.2/32 is directly connected, GigabitEthernet1/0/1
C        172.16.24.0/24 is directly connected, GigabitEthernet1/0/2
L        172.16.24.2/32 is directly connected, GigabitEthernet1/0/2
C        172.16.25.0/24 is directly connected, GigabitEthernet1/0/3
L        172.16.25.2/32 is directly connected, GigabitEthernet1/0/3
O        172.16.254.3/32
           [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
O        172.16.254.4/32
           [110/2] via 172.16.24.4, 03:37:28, GigabitEthernet1/0/2
O        172.16.254.5/32
           [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
O        172.16.255.1/32
           [110/3] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
           [110/3] via 172.16.24.4, 03:37:38, GigabitEthernet1/0/2
           [110/3] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
C        172.16.255.2/32 is directly connected, Loopback0
O        172.16.255.3/32
           [110/2] via 172.16.23.3, 05:36:24, GigabitEthernet1/0/1
O        172.16.255.4/32
           [110/2] via 172.16.24.4, 03:37:34, GigabitEthernet1/0/2
O        172.16.255.5/32
           [110/2] via 172.16.25.5, 03:39:11, GigabitEthernet1/0/3
Spine-02#
```

Return to Verify Dual-Homing with Single-Active Redundancy in a BGP EVPN VXLAN Fabric on page 26 .

Documents the configuration and verification of dual-homing with single-active redundancy in a BGP EVPN VXLAN fabric.

Part I

All-Active Mode

Topics:

This part contains the all-active mode, hierarchical multihoming, overlay: routed, DAG routed, and DAG bridged chapters.

2 All-Active Mode

Topics:

This chapter provides information about multihoming in all-active mode for fabric networks.

EVPN Multihoming Overview

Provides an overview of EVPN multihoming.

The Ethernet VPN (EVPN) multihoming is RFC 7432, RFC 8365 based industry-standard Layer 2 multipath solution solving various traditional networking protocol challenges impacting scale, performance and resiliency. It is an advanced networking technology designed to provide high availability, Border Gateway Protocol (BGP) based Layer 2 loop-free network, and efficient load balancing in modern enterprise campus networks.

A pair of Cisco Catalyst 9000 series switches can connect to downstream Layer 2 or Layer 3 network device with local physical connection binding into single logical EtherChannel configured in Layer 2 trunk or access mode. With industry-standard non-blocking architecture, any downstream networking devices such as, Ethernet Switch, Wireless LAN Controller (WLC), firewalls, servers and hosts can be dual-homed.

EVPN Multihoming Key Benefits

Explains the benefits of multihoming.

EVPN multihoming technology was introduced on Cisco Catalyst 9000 series switches in Cisco IOS XE 17.18.2, delivering a highly flexible and extensible deployment solution for global enterprise customers.

The key technology architecture benefits of Cisco Catalyst 9000 series switches combined with the advanced Cisco IOS XE software capabilities enable best of both worlds. Enterprise network administrators can implement EVPN multihoming within their existing traditional Layer 2 or Layer 3 network designs while laying a practical and gradual foundation for evolving towards modern, secure, fabric-based virtual networks.

The following illustration shows the EVPN multihoming solution for Cisco Catalyst 9000 series switches in two key flexible deployment models.

Figure 5: EVPN multihoming enterprise campus deployment model

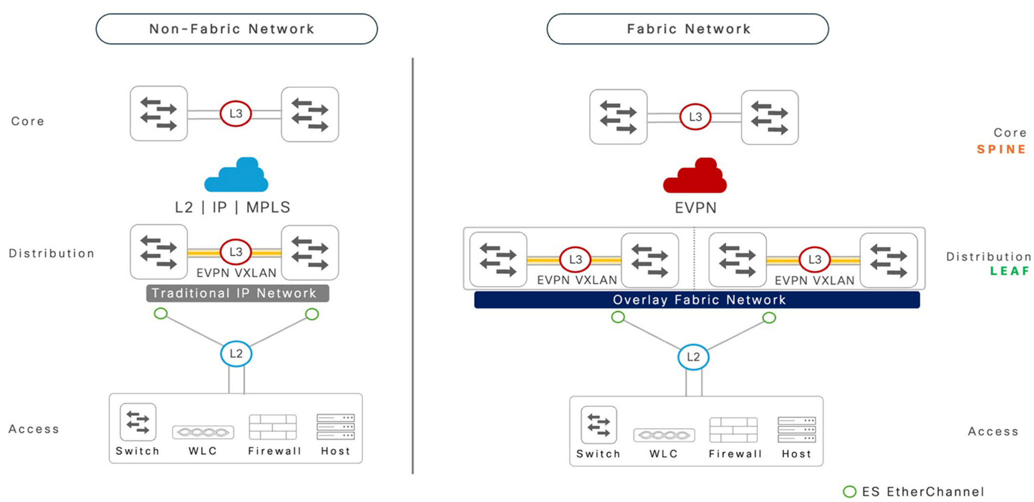

With flexible deployment model support, EVPN multihoming provides significant technological benefits to support enterprise-grade networking requirements.

- **Industry-standard Layer 2 multipath**: RFC 7432 and RFC 8365-defined Layer 2 multipath technology with BGP control-plane eliminates the vendor proprietary complex solutions.
- **Flexible architecture**: Purpose-built solution supports a wide range of networking use-cases for traditional Layer 2 and Layer 3, to modernized secure campus fabrics enabling network-wide, two-tier virtual network segmentation and extensions.

- **Distributed planes**: A fully distributed architecture where each system maintains its own scalable and resilient control plane, independent management plane configuration, and software versions provide more resilience with synchronized data-plane between paired systems.
- **Increased performance**: Non-blocking Spanning Tree Protocol (STP), Layer 2 loop-free network combined with advanced load-sharing technique assist in increasing forwarding capacity and improving application experience.
- **Seamless integration**: Least disruptive integration in classic Layer 2 networks with a wide range of networking device types that support link-bundling capabilities. Flexible and seamless integration with existing core network system without a forklift architecture change.
- **Resilient**: Deterministic fault detection and recovery technique during various types of planned and unplanned failure conditions, which reduces the Mean Time to Repair (MTTR) and increases the network reliability with a higher Mean Time Between Failure (MTBF).

  A pair of modular-class platforms with redundancy (Quad-Sup NSF and SSO): Provides non-stop business communication in maintaining network availability and forwarding throughput capacity during software upgrades and other unplanned failure events.

Cisco Catalyst 9000 Series Platform Support Matrix

Explains the platform support modes for EVPN multihoming.

EVPN multihoming is supported on a wide range of modular and fixed Cisco Catalyst 9000 series switch models. The system can be deployed in standalone or resilient mode, providing single network and device-level redundancy.

This table provides a list of supported Cisco Catalyst 9000 series platforms with the system mode.

Table 3: EVPN multihoming platform support matrix

	Cisco Catalyst 9000 series switch model	System mode
Modular platforms	• Catalyst 9400 Series Supervisor 1 • Catalyst 9400 Series Supervisor 2 • Catalyst 9400 Series Supervisor 2XL • Catalyst 9600 Supervisor 1	• Single supervisor: non-redundant • Dual supervisor: redundant
Fixed platforms	Catalyst 9300 Series	• Standalone: non-redundant • StackWise: redundant
	Catalyst 9500 High Performance Series	Standalone: non-redundant

Restrictions for EVPN Multihoming in All-Active Mode

Lists the unsupported platforms and describes the restrictions for EVPN multihoming in all-active mode.

This section lists the deployment restrictions for multihoming in all-active mode.

- EVPN multihoming is not supported on Cisco Catalyst 9500-X series switches, Catalyst 9600 Supervisor 2, Cisco C9350 series smart switches, and Cisco C9610 series smart switches.

- Combined mode with StackWise Virtual and EVPN multihoming on single switch is not supported on Cisco Catalyst 9400 series switches, Cisco Catalyst 9500-High Performance series switches, and Cisco Catalyst 9600 series switches.
- EVPN multihoming for non-fabric and fabric networks with IPv6-only underlay network is not supported.
- EVPN multihoming is not supported with Layer 2 leaf mode in centralized gateway mode fabric deployments.
- EVPN multihoming In-Service Software Upgrade (ISSU) from a release prior to Cisco IOS XE 17.18.2 is not supported on Cisco Catalyst 9400 series and Cisco Catalyst 9600 series modular platforms with redundant supervisors.
- Anycast gateway (global or VRF-enabled) SVI with First Hop Redundancy Protocol (FHRP) protocol is not supported.
- Resilient Ethernet Protocol (REP) on Layer 2 Ethernet segment ports is not supported.
- Private VLAN (PVLAN) on Layer 2 Ethernet segment ports is not supported.
- DHCP snooping and ARP inspection on VLAN IDs mapped on logical Ethernet Segment EtherChannel trunk port is not supported.

EVPN Scale Matrix in Fabric Mode

Provides system-wide layer 2 interface and table entry scale numbers for EVPN multihoming in fabric mode.

This section describes the system-wide Layer 2 interface and table entry scale limits for EVPN multihoming enabled on Cisco Catalyst 9000 series switches in fabric mode.

Table 4: EVPN multihoming platform scale matrix

EVPN multihoming segment scale matrix	Scale count
Ethernet segment switch per redundancy group	2
Ethernet segment Port Channel interface	48
Number of VLANs	200
MAC address	10,000
IPv4 address	10,000
IPv6 address	20,000
IP VRF	25

EVPN Multihoming Technology Overview

Provides a brief overview of the industry-standard EVPN multhoming technology for traditional Layer 2 and Layer 3 networks.

The legacy networking protocols in enterprise campus networks have been a challenge for IT organizations.

The challenges include difficulty in eliminating Spanning Tree Protocol (STP) that can lead to inefficient network topologies, restrictions in traditional Layer 2 designs that limit the network switching capacity, lack of deterministic reliability that makes it a challenge to support real-time mission-critical applications, insufficient support for mobility in wireless and legacy application environments, and increased complexity in network management and troubleshooting.

This section provides a brief overview of the advanced industry-standard EVPN multihoming technology for traditional Layer 2 and Layer 3 networks. Each subsection focuses on various technology components enabling loop-free, high-performance and resilient Layer 2 networking solutions.

The following illustration captures the key benefits of EVPN multihoming technology for simplified and resilient enterprise campus networks.

Figure 6: EVPN multihoming technology benefits

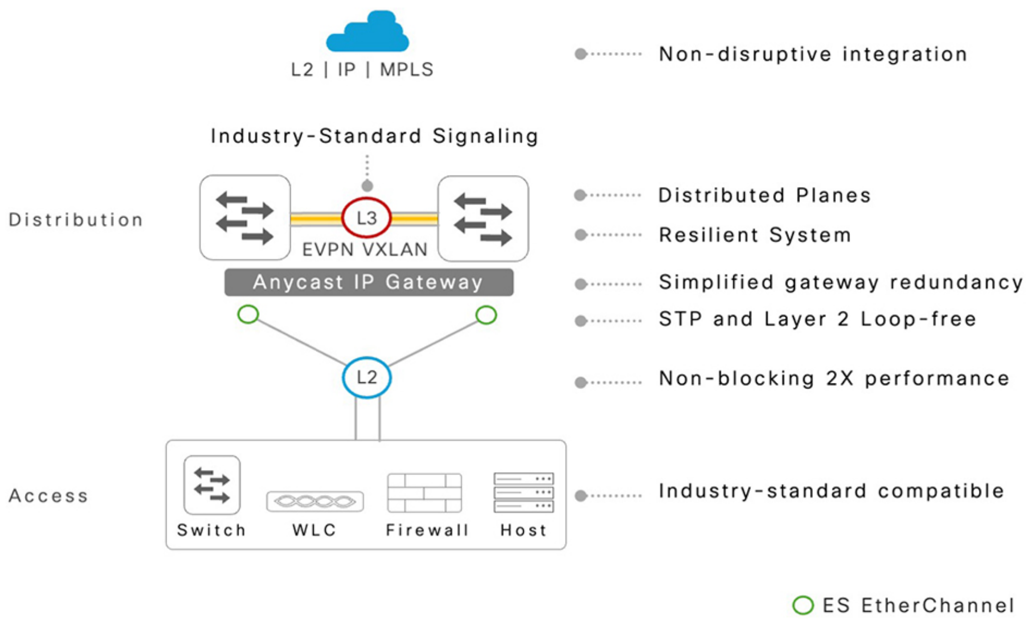

Refer to the Terms and Definitions section for the description and purpose of acronyms.

Network System Components

Describes the various system components of EVPN multi-homing in non-fabric networks.

The EVPN multihoming technology comprises directly attached networking systems through traditional Layer 2 ports in trunk or access mode. The Cisco Catalyst 9000 series switches are required to provision in EVPN multihoming mode to provide a loop-free Layer 2 multipath EtherChannel to any type of downstream network device.

Ethernet Segment Redundancy Group

An Ethernet segment redundancy group is a physical grouping of network devices that form an EVPN multihoming network, enabling an all-active Layer 2 multipath solution.

The EVPN multihoming redundancy group consists of a single pair of Cisco Catalyst 9000 series switches, which provide physical connectivity and system-level redundancy. This architecture ensures continuous communication during both planned and unplanned failure events.

Ethernet Segment Switch

An Ethernet segment switch is a typical Cisco Catalyst 9000 series switch in core or distribution system configured to support the Layer 2 multihomed network with directly attached various types of Layer 2 systems.

The Ethernet Segment (ES) system can support multifunctional Layer 2 and Layer 3 networks for enterprise campus networks in classic and traditional non-fabric or modern fabric network roles. As a result, the ES switch is also known as a Provider Edge (PE) or Leaf or VTEP. Each ES switch operates and maintains fully distributed control, management, and data planes while solving legacy protocol challenges with industry-standard Layer 2 multipath solutions.

Note:

Cisco Catalyst 9000 series switches in StackWise mode are considered as logical system Ethernet segment switches.

Ethernet Segment Switch Platform and Software

The EVPN multihoming technology operation between ES network devices–fixed or modular models, is no different from any of the other industry-standard networking protocols, like OSPF and BGP. A pair of ES switches can have different platform types, modules, interfaces, and so on. The ES pair switches can also have different Cisco IOS XE software versions to address stage upgrades and other conditions.

While the technology permits asymmetric platform and software versions on an ES pair, Cisco recommends a common platform and software version to deliver consistent performance and resiliency for non-disruptive business continuity during planned and unplanned failure conditions.

Ethernet Segment Client

The Layer 2 system is a directly attached single, dual, or multihomed connection to a pair of Cisco Catalyst 9000 series switches in an ES system. The Layer 2 ES client network can be of any device type but must follow the industry-standard Layer 2 networking technologies.

EVPN Multihoming Control Plane

Describes the uses of BGP control-plane in multi-homing.

The EVPN multihoming technology is built on industry-standard, highly flexible, and proven BGP routing protocol. The BGP control plane replaces legacy STP for Layer 2 loop detection and prevention techniques.

The L2VPN EVPN address-family is a multiprotocol extension that enables a network-agnostic multihoming solution. The BGP protocol provides four key functions: discover remote neighbor and ES ID, real-time synchronization of network states, distributed forwarding rule to select local-bias rules for optimal performance, and resiliency for deterministic and efficient rapid fault-detection and recovery.

Ethernet Segment

Describes how to enable an Ethernet segment in the aggregation layer.

Ethernet segment (ES) is a pair of Cisco Catalyst 9000 series switches in the aggregation layer with a Layer 2 physical port that directly connects to an ES client device. It represents a single logical entity to enable loop-free, non-blocking all-active Layer 2 network connectivity.

An STP-Bridge Protocol Data Unit (BPDU) free Layer 2 network dynamically supports per-VLAN loop prevention while maximizing the network throughput to support accelerated application performance and resiliency.

Ethernet Segment ID

Explains Ethernet Segment ID in the context of Cisco Catalyst 9000 series switches and EVPN multi-homing.

An Ethernet segment ID is a 10-byte (00:01:01:01:01:01:01:01:01:01) identifier for each Layer 2 port that is connected to an ES client device. A common ES ID is assigned to a pair of Catalyst 9000 ES systems on the Layer 2 physical port that connects to the same Layer 2 ES client device to enable all-active EVPN multihoming.

Cisco Catalyst 9000 series switches support 3 types of industry-standard ESI ID that is either auto-generated or IT-defined and manually configured.

Anycast Gateway MAC and IP Addresses

How anycast gateway MAC and IP address works in an EVPN multi-homing environment

A Cisco Catalyst 9000 SVI interface with anycast gateway is provisioned with a shared virtual IP address and an auto-derived shared virtual MAC address with **anycast-gateway mac auto** command in the global settings, on pair of EVPN multihoming switches. The unified, resolved ARP and ND addresses between the distributed Catalyst 9000 IP gateway switches support optimal data load-balancing across all available bundled links and resilient networks during planned or unplanned failures for Wired and Wireless endpoints.

Designated Forwarder and Non-Designated Forwarder Roles

Describes the roles of designated forwarder and non-designated forwarder.

Traditionally, in Layer 2 STP-enabled networks, loop detection is achieved through a blocking link to transmit broadcast and business application data traffic. As a result, the network operates inefficiently at reduced bandwidth capacity. Additionally, protocol-based fault detection and recovery increase the network convergence time, which impacts the reliability of mission-critical applications during faults.

In EVPN multihoming-based campus networks, the same result is obtained using a very different logic. Instead of blocking links, the EVPN distinctly decouples network traffic types between broadcast categories and business applications. On a per ES port and VLAN basis, the pair of Cisco Catalyst 9000 series switches dynamically assign specific roles to forward BUM (Broadcast, Unknown Unicast, and Multicast) traffic in active and passive modes. However, business application forwarding traffic is unblocked in active-active mode.

The Catalyst 9000 series switch elected to actively forward the BUM traffic on a shared ES EtherChannel interface is known as Designated Forwarder (DF). The Catalyst 9000 series switch elected to block the forwarding of the BUM traffic on a shared ES EtherChannel interface is known as Non-Designated Forwarder (non-DF). The dynamic DF and non-DF role assignment is automatically derived based on the internal system modulo hash algorithm, which enables auto load balancing of the BUM traffic between a pair of Catalyst 9000 switches connected to same downstream Layer 2 network access devices.

Figure 7: Designated forwarder and non-designated forwarder roles

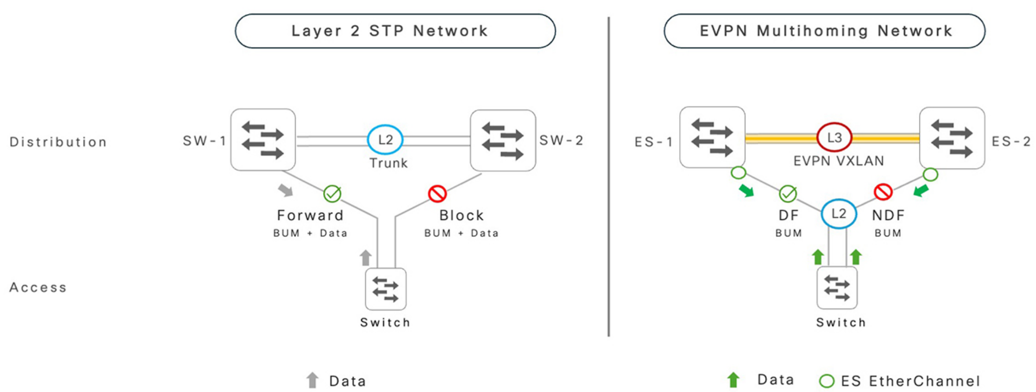

Dynamic MAC and IP Learning and Synchronization

How dynamic MAC and IP address learning and synchronization works in an ESI multihoming environment.

EVPN multihoming networks build and maintain the Layer 2 and Layer 3 network information using the control plane, between a pair of Cisco Catalyst 9000 ES switches.

The upstream data towards the IP core network is hashed from Layer 2 network devices, the dynamic MAC, IPv4, and IPv6 host addresses synchronized using the BGP control plane in real time between both the ES switches.

The common MAC or IP forwarding tables enable high performance and fully-distributed local forwarding while pre-programming the inter-ES Layer 2 VXLAN tunnel to bridge the downstream Layer 2 network traffic as the last-resort interface to reroute the data plane rapidly, upon local path failure, without relying on data-plane flooding.

IGMP Join and Leave Synchronization

Explains how IGMP join andl eave messages are synchronized across multihomed Ethernet switches.

The incoming Internet Group Management Protocol (IGMP) messages from multicast host receivers are locally processed by the connected Cisco Catalyst 9000 ES switches. The IPv4 or IPv6 multicast group-to-IP membership information is synchronized between the ES switches by using an extended BGP control plane that supports consistent multicast state across the multihomed Ethernet segments.

The Catalyst 9000 ES switch with a VLAN in the DF role transmits egress multicast traffic towards the receiver. The peer non-DF ES switch suppresses duplicate multicast frames to prevent loops and undesired multicast replication. This enables EVPN multihoming to support symmetric unicast or multicast application performance and resiliency during planned or unplanned failure events.

EVPN Multihoming for Non-Fabric Networks

An overview of EVPN multihoming for non-fabric networks.

EVPN multihoming in Cisco Catalyst 9000 series switches provides a flexible deployment solution for global enterprise customers to retain their traditional Layer 2 or Layer 3 networks in upstream core networks. Such network deployment option is also known as non-fabric network.

EVPN Multihoming in Fabric Networks

Explains how EVPN multihoming works in fabric networks.

Modern enterprise campuses require a secure fabric networking solution that seamlessly supports large-scale virtual and logical networks over a robust physical network infrastructure. EVPN multihoming-based networks offer a simplified and resilient architecture evolving from legacy STP protocol-based enterprise campus networks to a non-blocking Layer 2 fabric.

The EVPN VXLAN-enabled fabric core networks introduce new possibilities for EVPN multihomed Layer 2 networks by enabling secured wired and wireless segmented and extended networks that address critical technical and business requirements.

The following illustration displays EVPN multihoming for fabric networks

Figure 8: EVPN multihoming for fabric networks

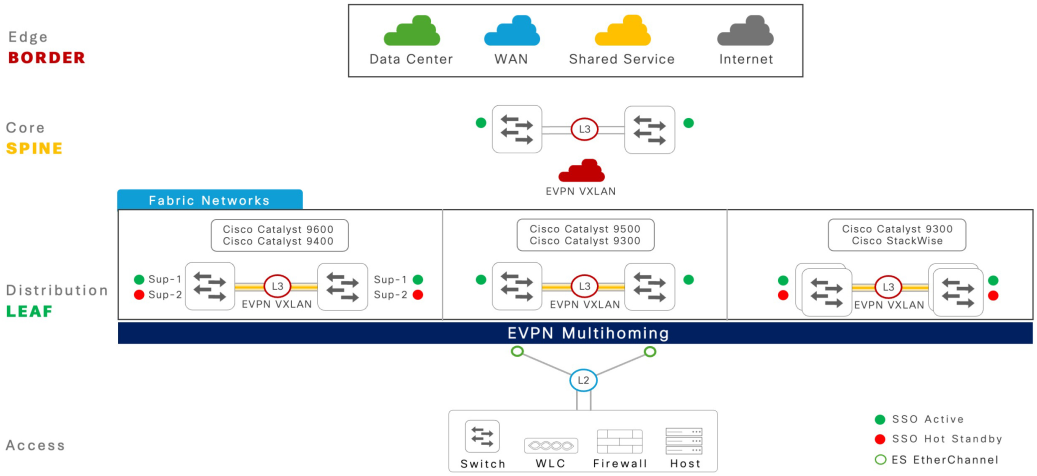

Single Unified BGP Control Plane

Explains how a single unified BGP control plane streamlines enterprise campus network operations by integrating traditional Layer 2 and Layer 3 access with modern fabric core networking.

Enterprise campuses can be designed and deployed using a single unified BGP control plane that addresses the traditional Layer 2 and Layer 3 networking requirements at the access layer while supporting a modern fabric network at the core. The unified BGP control plane simplifies network operations by distinctly managing multi-domain routing and bridging functions, enabling seamless integration of the access and core networks.

BGP EVPN multihoming-based networks represents a shift from the traditional flood-based networks to control-plane driven architecture that supports non-blocking and resilient enterprise campus networks. Host addresses–including the MAC, IPv4 or IPv6, multicast–are statefully discovered and synchronized across peer systems, enabling all-active data forwarding in BGP EVPN multihoming networks.

As the BGP control plane role expands into traditional networks, processing demands may increase due to additional BGP prefix tables supporting EVPN multihoming auto-generated prefixes.

Cisco Catalyst 9000 series switches support both standard non-hierarchical and hierarchical BGP peering mode between leaf and spine devices, facilitating the exchange of fabric and EVPN multihoming network prefixes.

To address scalability and performance requirements, Catalyst 9000 series switches can implement a two-tier hierarchical control-plane network in large deployments.

Figure 9: BGP peering alternatives for fabric networks

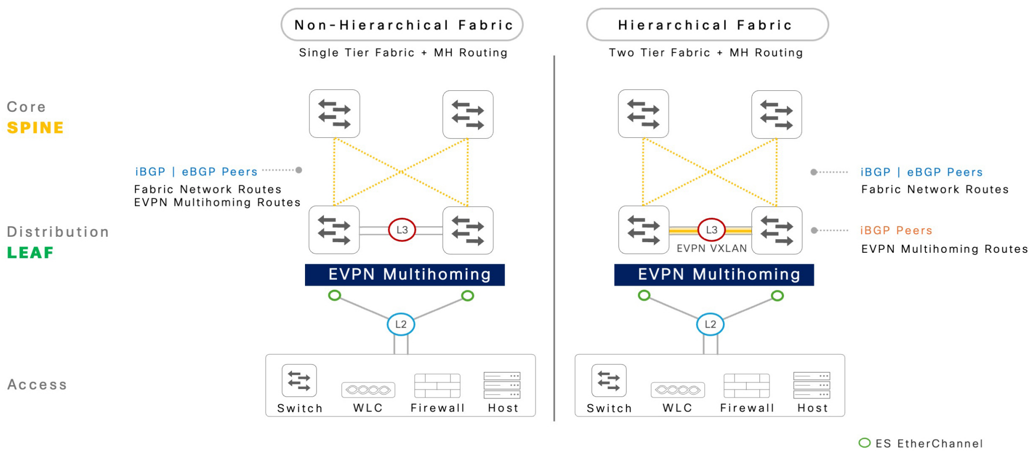

Non-Hierarchical Fabric Networks

Explains how non-hierarchical fabric networks operate in large-scale BGP EVPN VXLAN implementations, route exchanges, and the resource requirements for spine devices.

The general non-hierarchical BGP EVPN VXLAN network implementation follows a standard single tier BGP peering model where all fabric devices peer directly with each other to dynamically exchange routes and build the VXLAN forwarding tunnels.

The EVPN multihoming-enabled non-hierarchical fabric network implementation follows a standard iBGP or eBGP peering model between the spine and EVPN multihoming-enabled leaf or border system.

In large scale enterprise campus network environments, spine layer devices require additional resources such as processing power and memory to handle standard fabric host information, such as MAC addresses, MAC/IP bindings and network prefixes (IPv4/IPv6) along with EVPN multihoming auto-generated routes received from each leaf or border system.

Hierarchical Fabric Networks

Explains how hierarchical fabric networks use structured networking and BGP-based designs to provide scalable and resilient campus network architectures.

Hierarchical and structured networking is fundamental to supporting scalable and resilient campus networks, and EVPN multihoming-based fabric networks follow hierarchical BGP routing design principles for better scalable networking solutions.

The two-tier BGP peering in EVPN multihoming networks assists in subdividing route management between peers by separating Layer 2 networking from scalable network connectivity towards the spine layer.

The iBGP peering between a pair of Cisco Catalyst 9000 series switches exchanges all auto-generated EVPN multihoming network prefixes to build reliable and scalable traditional Layer 2 networks.

Cisco Catalyst 9000 series switches in EVPN multihoming leaf role follow standard parallel iBGP or eBGP peering with a pair of spine switches advertising network prefixes based on EVPN fabric overlay network types. The conditional network prefix announcement to spine switches enables enhanced flexibility, scale, performance, and resiliency across the network.

For more information on overlay networks, refer to Hierarchical EVPN Control-Plane.

EVPN Fabric Overlay Network Types Overview

Explains and compares the different overlay network types in EVPN fabric multihoming deployments.

BGP EVPN VXLAN provides a flexible, virtual networking solution that addresses key networking use cases by combining a wide range of overlay architectures.

Such business-driven network architectures enable multidimensional benefits for scalable and secure segmented overlay networks in EVPN multihoming networks. Network administrators can build overlay networks with a “route first” mindset that can conditionally extend IP or VLAN networks between targeted Ethernet segment network devices to meet application requirements.

The following figure illustrates three overlay network types in EVPN multihoming networks.

Figure 10: BGP EVPN fabric overlay network types

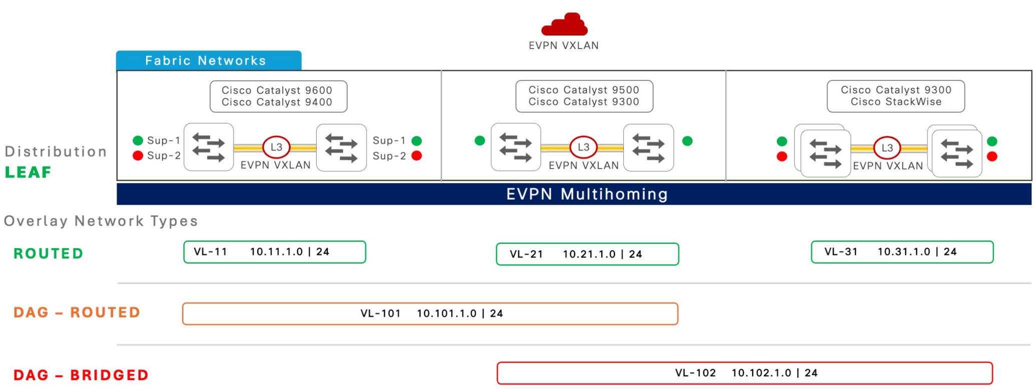

Network administrators decide on the type of overlay networks to deploy across the fabric core based on the specific application requirements to achieve better scale and resiliency in EVPN multihoming networks.

Cisco Catalyst 9000 series switches provide a flexible solution that supports coexistence of all overlay network types within a single system. Based on technical requirements, each VLAN from an Ethernet segment EtherChannel on Cisco Catalyst 9000 series switches can be configured to support IP-routed network segmentation, conditionally stretch IP subnets, or bridge VLANs using Layer 2 flood mechanisms selectively between VLANs.

Table 5: Comparison of EVPN overlay network types

	EVPN overlay network types		
	Routed	**DAG-routed**	**DAG-bridged**
Function	IP routing in core and Layer 3 segmented overlay.	Flood-free IP subnet stretch and Layer 3 segmented overlay.	Layer 2 flood stretch and Layer 3 segmented overlay.
Use case	Wired or wireless data. IT or OT endpoints	Wired or wireless data. IT or OT endpoints	Wired or wireless data. IT or OT endpoints. Non-IP end-points
VLAN/subnet	One VLAN or subnet per dis-tribution block	Stretched IP subnets be-tween targeted distribution blocks	Stretched VLANs or subnets between targeted distribu-tion blocks

	EVPN overlay network types		
	Routed	**DAG-routed**	**DAG-bridged**
Layer 2 flood boundary	Within the local Layer 2 network	Within the local Layer 2 network	Within the local Layer 2 network and across the fabric core
IP gateway	Anycast gateway per distribution block	Anycast gateway between targeted distribution blocks	Anycast gateway between targeted distribution blocks
IP subnet stretch	Not applicable	Yes–flood-free	Yes–flood-based
Seamless distributed wireless mobility	Within the local distribution block	Within the local distribution block and across targeted destination blocks	Within the local distribution block and across targeted destination blocks
Non-IP/silent host support	Within the local distribution block	Within the local distribution block	Within the local distribution block and across targeted destination blocks
Recommendation	Best scalable solution	Extends IP subnet selectively; use only if required	Extends Layer 2 flood selectively; use only if required

EVPN Fabric and Non-Fabric Network Co-Existence

Explains how Cisco Catalyst 9000 series switches enable the coexistence of modern EVPN fabric networks and traditional Layer 2 or Layer 3 networks on the same campus infrastructure.

Enterprise campuses require flexible networking solutions that support both modern secure fabric networks and backward-compatible traditional Layer 2 or Layer 3 networks. Cisco Catalyst 9000 series switches with EVPN multihoming-enabled EtherChannel trunk interfaces provide this flexibility by bundling multiple VLANs, each delivering distinct networking services.

Network administrators can program some VLANs to operate traditional IP-based networks in the underlay, while mapping other VLANs to IP VRFs or MAC VRFs to route and bridge over the VXLAN fabric.

The following illustration shows a network deployment scenario with both fabric and non-fabric VLANs on EVPN multihoming-enabled networks.

Figure 11: EVPN multihoming: fabric and non-fabric network co-existence

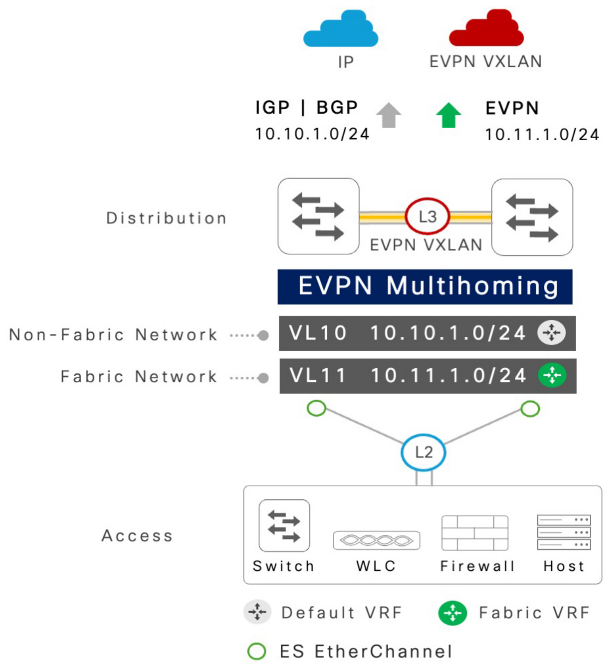

To understand and implement EVPN multihoming for non-fabric deployments, refer to