



L2VPN and Ethernet Services Configuration Guide for Cisco NCS 540 Series Routers, IOS XR Release 6.3.x

First Published: 2017-03-30

Americas Headquarters

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

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

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

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

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

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

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

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

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

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

© 2018 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface ix

Obtaining Documentation and Submitting a Service Request ix

CHAPTER 1

Configure Gigabit Ethernet for Layer 2 VPNs 1

Introduction to Layer 2 Virtual Private Networks 1

Introduction to Layer 2 VPNs on Gigabit Ethernet Interfaces 1

Configure Gigabit Ethernet Interfaces for Layer 2 Transport 3

Ethernet Data Plane Loopback 4

Configure Ethernet Data Plane Loopback 6

Running Configuration 7

Verification 8

CHAPTER 2

L2CP Tunneling MEF 11

L2CP Tunneling 11

Configure L2CP Tunneling 12

CHAPTER 3

Configure Layer 2 Access Control Lists 13

Layer 2 Access Control Lists 13

Prerequisites for Configuring Layer 2 Access Control Lists 13

Layer 2 Access Control Lists Feature Highlights 14

Purpose of Layer 2 Access Control Lists 14

How a Layer 2 Access Control List Works 14

Layer 2 Access Control List Process and Rules 14

Create Layer 2 Access Control List 15

Restrictions for Configuring Layer 2 Access Control Lists 15

Configuration 15

Running Configuration 16
 Verification 16

CHAPTER 4

Configure Link Bundles for Layer 2 VPNs 17
 Configure Gigabit Ethernet Link Bundle 17
 Configure VLAN Bundle 20
 References for Configuring Link Bundles 21
 Characteristics of Link Bundles 21
 Methods of Forming Bundles of Ethernet Interfaces 22
 Link Aggregation Through LACP 23

CHAPTER 5

Configure Point-to-Point Layer 2 Services 25
 Ethernet over MPLS 25
 Ethernet Port Mode 26
 VLAN Mode 26
 QinQ Mode 27
 Configure Local Switching Between Attachment Circuits 28
 Configure Static Point-to-Point Connections Using Cross-Connect Circuits 32
 Flexible Cross-Connect Service 34
 Flexible Cross-Connect Service - Single-Homed 34
 Flexible Cross-Connect Service - Multi-Homed 34
 Flexible Cross-Connect Service Supported Modes 35
 VLAN Unaware 35
 Configure Single-Homed Flexible Cross-Connect Service using VLAN Unaware 36
 Configure Multi-Homed Flexible Cross-Connect Service using VLAN Unaware 37
 VLAN Aware 41
 Configure Single-Homed Flexible Cross-Connect using VLAN Aware 41
 Configure Multi-Homed Flexible Cross-Connect Service using VLAN Aware 43
 Local Switching 47
 Configure Multi-Homed Flexible Cross-Connect Service using Local Switching 47
 Configure Preferred Tunnel Path 49
 Configure Pseudowire Redundancy 50
 Access Pseudowire Redundancy 51
 Configure Access Pseudowire Redundancy 51

Virtual Circuit Connection Verification on L2VPN	53
GTP Load Balancing	53

CHAPTER 6**Configure Virtual LANs in Layer 2 VPNs 57**

Configure VLAN Sub-Interfaces	58
Introduction to Ethernet Flow Point	60
Identify Frames of an EFP	61
Apply Features	61
Define Data-Forwarding Behavior	62
Ethernet Flow Points Visibility	62
Configuring EFP Visibility	63
Configure VLAN Header Rewrite	65
Valid Ingress Rewrite Actions	68
Valid Ingress-Egress Rewrite Combinations	69

CHAPTER 7**EVPN Features 79**

EVPN Overview	79
EVPN Concepts	80
EVPN Operation	81
EVPN Route Types	82
Configure EVPN L2 Bridging Service	83
Running Configuration	84
EVPN Software MAC Learning	84
Configure EVPN Software MAC Learning	85
Supported Modes for EVPN Software MAC Learning	85
Single Home Device or Single Home Network Mode	86
Configure EVPN in Single Home Device or Single Home Network Mode	86
Dual Home Device—All-Active Load Balancing Mode	87
Configure EVPN Software MAC Learning in Dual Home Device—All-Active Mode	88
Verify EVPN Software MAC Learning	90
EVPN Out of Service	92
Configure EVPN Out of Service	93
Running Configuration	93
EVPN Multiple Services per Ethernet Segment	95

Configure EVPN Multiple Services per Ethernet Segment	96
Configuration Example	96
Running Configuration	98
Associated Commands	101
Network Convergence using Core Isolation Protection	101
Configure EVPN Convergence using Core Isolation Protection	103
Conditional Advertisement of Default-Originate	107
Configure Conditional Advertisement of Default-Originate	108
EVPN Routing Policy	111
EVPN Route Types	111
EVPN RPL Attribute	116
EVPN RPL Attribute Set	118
Configure EVPN RPL Feature	119
Running Configuration	120
Support for DHCPv4 and DHCPv6 Client over BVI	126
Configure DHCPv4 and DHCPv6 Client over BVI	126

CHAPTER 8**Configure EVPN IRB 133**

EVPN IRB	133
EVPN Single-Homing Access Gateway	134
EVPN Multihoming All-Active	135
Enable Auto-BGP RT with Manual ESI Configuration	136
Supported EVPN IRB Scenarios	136
Distributed Anycast Gateway	136
EVPN IRB with All-Active Multi-Homing without Subnet Stretch or Host-Routing across the Fabric	137
EVPN IRB with All-Active Multihoming with Subnet Stretch or Host-Routing across the Fabric	137
MAC and IP Unicast Control Plane	138
Intra-subnet Unicast Data Plane	139
Inter-subnet Unicast Data Plane	139
VM Mobility Support	139
MAC and MAC-IP Sequence Numbers	140
Synchronized MAC and MAC-IP Sequence Numbers	140

Local Sequence Number Updates	140
Best Route Selection after Host Movement	140
Stale Route Deletion after a Host Movement	140
Host Movement Detection through GARP	141
Host Move Detection with Silent Host	141
Host Move Detection without GARP with Data Packet	141
Duplicate MAC Detection	141
Configuring EVPN IRB	141
Running Configuration for EVPN IRB	143
Verify EVPN IRB	144
EVPN IPv6 Hosts with Mobility	144
Configure EVPN IPv6 Hosts with Mobility	145
Running Configuration	149
Verification	153
<hr/>	
CHAPTER 9	EVPN Virtual Private Wire Service (VPWS) 157
EVPN-VPWS Single Homed	157
Configure EVPN-VPWS Single Homed	158
Running Configuration	158
EVPN-VPWS Multi-Homed	159
Configure EVPN-VPWS Multi-Homed	159
Running Configuration	161
<hr/>	
CHAPTER 10	L2VPN Services over Segment Routing for Traffic Engineering Policy 163
EVPN VPWS Preferred Path over SR-TE Policy	164
Configure EVPN VPWS Preferred Path over SR-TE Policy	165
Configure Prefix-SID in ISIS	165
Configure Adjacency-SID in ISIS	167
Configure Segment-list	169
Configure SR-TE Policy	169
Configure EVPN VPWS over SR-TE Policy	170
Running Configuration	171
Verify EVPN VPWS Preferred Path over SR-TE Policy Configuration	175
Associated Commands	176

- Related Topics 176
- L2VPN VPWS Preferred Path over SR-TE Policy 176
 - Configure L2VPN VPWS Preferred Path over SR-TE Policy 176
 - Configure Prefix-SID in IS-IS 177
 - Configure Adjacency-SID in IS-IS 178
 - Configure Segment-list 180
 - Configure SR-TE Policy 181
 - Configure VPWS over SR-TE Policy 182
 - Running Configuration 182
 - Verify L2VPN VPWS Preferred Path over SR-TE Policy Configuration 186
 - Associated Commands 188
 - Related Topics 188
 - EVPN VPWS On-Demand Next Hop with SR-TE 189
 - Configure EVPN VPWS On Demand Next Hop with SR-TE 189
 - Topology 190
 - Configure Prefix-SID in ISIS 190
 - Configure SR-TE 192
 - Configure PCE and PCC 193
 - Configure SR Color 193
 - Configure EVPN Route Policy 194
 - Configure BGP 194
 - Configure EVPN VPWS 195
 - Configure Flexible Cross-connect Service (FXC) VLAN-unaware 196
 - Running Configuration 196
 - Related Topics 203
 - Overview of Segment Routing 203
 - How Segment Routing Works 204
 - Segment Routing Global Block 205

CHAPTER 11

MSTP BPDU Guard 207

- Configuring MSTP BPDU Guard 207
 - Running Configuration with MSTP BPDU Guard 208
 - Verification for MSTP BPDU Guard 208



Preface



Note This product has reached end-of-life status. For more information, see the [End-of-Life and End-of-Sale Notices](#).

This preface contains these sections:

- [Obtaining Documentation and Submitting a Service Request](#), on page ix

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/en/US/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.



CHAPTER 1

Configure Gigabit Ethernet for Layer 2 VPNs

This chapter introduces you to Layer 2 features and standards, and describes how you can configure L2VPN features.

The distributed Gigabit Ethernet (including 10-Gigabit and 100-Gigabit) architecture and features deliver network scalability and performance, while enabling service providers to offer high-density, high-bandwidth networking solutions designed to interconnect the router with other systems in POPs, including core and edge routers and Layer 2 and Layer 3 switches.

- [Introduction to Layer 2 Virtual Private Networks, on page 1](#)
- [Introduction to Layer 2 VPNs on Gigabit Ethernet Interfaces, on page 1](#)
- [Configure Gigabit Ethernet Interfaces for Layer 2 Transport, on page 3](#)
- [Ethernet Data Plane Loopback, on page 4](#)

Introduction to Layer 2 Virtual Private Networks

A Layer 2 Virtual Private Network (VPN) emulates a physical sub-network in an IP or MPLS network, by creating private connections between two points. Building a L2VPN network requires coordination between the service provider and customer. The service provider establishes Layer 2 connectivity. The customer builds a network by using the data link resources obtained from the service provider. In a L2VPN service, the service provider does not require information about the customer's network topology and other information. This helps maintain customer privacy, while using the service provider resources to establish the network.

The service provider requires Provider Edge (PE) routers with the following capabilities:

- Encapsulation of L2 protocol data units (PDU) into Layer 3 (L3) packets.
- Interconnection of any-to-any L2 transports.
- Support for MPLS tunneling mechanism.
- Process databases that include all information related to circuits and their connections.

This section introduces Layer 2 Virtual Private Networks (VPNs) and the corresponding Gigabit Ethernet services.

Introduction to Layer 2 VPNs on Gigabit Ethernet Interfaces

A L2VPN network enables service providers (SPs) to provide L2 services to geographically disparate customer sites. Typically, a SP uses an access network to connect the customer to the core network. This access network may use a mixture of L2 technologies, such as Ethernet and Frame Relay. The connection between the customer

site and the nearby SP edge router is known as an attachment circuit (AC). Traffic from the customer travels over this link to the edge of the SP core network. The traffic then tunnels through a pseudowire over the SP core network to another edge router. The edge router sends the traffic down another AC to the customer's remote site.

The L2VPN feature enables the connection between different types of L2 attachment circuits and pseudowires, allowing users to implement different types of end-to-end services.



Note BOOTP traffic (dst UDP 68) over any type of pseudowire is unsupported.

Cisco IOS XR software supports a point-to-point end-to-end service, where two Ethernet circuits are connected together. An L2VPN Ethernet port can operate in one of two modes:

- **Port Mode**—In this mode, all packets reaching the port are sent over the pseudowire, regardless of any VLAN tags that are present on the packets. In Port mode, the configuration is performed under the `l2transport` configuration mode.
- **VLAN Mode**—Each VLAN on a CE (customer edge) or access network to PE (provider edge) link can be configured as a separate L2VPN connection (using either VC type 4 or VC type 5). To configure L2VPN on VLANs, see *The Carrier Ethernet Model* chapter in this manual. In VLAN mode, the configuration is performed under the individual sub-interface.

Switching can take place in the following ways:

- **AC-to-PW**—Traffic reaching the PE is tunneled over a PW (pseudowire) (and conversely, traffic arriving over the PW is sent out over the AC). This is the most common scenario.
- **Local switching**—Traffic arriving on one AC is immediately sent out of another AC without passing through a pseudowire.
- **PW stitching**—Traffic arriving on a PW is not sent to an AC, but is sent back into the core over another PW.

**Note**

- If your network requires that packets are transported transparently, you may need to modify the packet's destination MAC (Media Access Control) address at the edge of the Service Provider (SP) network. This prevents the packet from being consumed by the devices in the SP network.
- The **encapsulation dot1ad *vlan-id*** and **encapsulation dot1ad *vlan-id* dot1q any** commands cannot co-exist on the same physical interface or bundle interface. Similarly, the **encapsulation dot1q *vlan-id*** and **encap dot1q *vlan-id* second-dot1q any** commands cannot co-exist on the same physical interface or bundle interface. If there is a need to co-exist, it is recommended to use the exact keyword in the single tag encapsulation. For example, **encap dot1ad *vlan-id* exact** or **encap dot1q *vlan-id* exact**.
- In an interface which already has QinQ configuration, you cannot configure the QinQ Range sub-interface where outer VLAN range of QinQ Range overlaps with outer VLAN of QinQ. Attempting this configuration results in the splitting of the existing QinQ and QinQ Range interfaces. However, the system can be recovered by deleting a recently configured QinQ Range interface.
- In an interface which already has QinQ Range configuration, you cannot configure the QinQ Range sub-interface where outer VLAN range of QinQ Range overlaps with inner VLAN of QinQ Range. Attempting this configuration results in the splitting of the existing QinQ and QinQ Range interfaces. However, the system can be recovered by deleting a recently configured QinQ Range interface.

You can use the **show interfaces** command to display AC and pseudowire information.

Configure Gigabit Ethernet Interfaces for Layer 2 Transport

This section describes how you can configure Gigabit ethernet interfaces for Layer 2 transport.

Configuration Example

```
RP/0/RP0/CPU0(config)#interface TenGigE 0/0/0/10

/* Configure the ethertype for the 802.1q encapsulation (optional) */
/* For VLANs, the default ethertype is 0x8100. In this example, we configure a value of
0x9100.
/* The other assignable value is 0x9200 */
/* When ethertype is configured on a physical interface, it is applied to all sub-interfaces
created on this interface */

RP/0/RP0/CPU0:router(config-if)#dot1q tunneling ethertype 0x9100

/* Configure Layer 2 transport on the interface, and commit your configuration */
RP/0/RP0/CPU0:router(config-if)#l2transport
RP/0/RP0/CPU0:router(config-if-l2)#commit
Sat May  2 19:50:36.799 UTC
RP/0/RP0/CPU0:router(config-if-l2)#exit
RP/0/RP0/CPU0:router(config-if)#no shutdown
RP/0/RP0/CPU0:router(config-if)#exit
RP/0/RP0/CPU0:router(config)#
```

Running Configuration

```
configure
interface TenGigE 0/0/0/10
```

```
dot1q tunneling ethertype 0x9100
l2transport
!
```

Verification

Verify that the Ten-Gigabit Ethernet interface is up and operational.

```
router# show interfaces TenGigE 0/0/0/10
...
TenGigE0/0/0/10 is up, line protocol is up
Interface state transitions: 1
Hardware is TenGigE, address is 0011.1aac.a05a (bia 0011.1aac.a05a)
Layer 1 Transport Mode is LAN
Layer 2 Transport Mode
MTU 1514 bytes, BW 10000000 Kbit (Max: 10000000 Kbit)
  reliability 255/255, txload 0/255, rxload 0/255
Encapsulation ARPA,
Full-duplex, 10000Mb/s, link type is force-up
output flow control is off, input flow control is off
Carrier delay (up) is 10 msec
loopback not set,
...
```

Associated Commands

- [l2transport \(Ethernet\)](#)

Ethernet Data Plane Loopback

The Ethernet Data Plane Loopback function allows you to run loopback tests to test the connectivity and quality of connections through a Layer 2 cloud. You can run this test on:

- Main interface or sub-interfaces
- Bundle or its sub-interfaces
- Multiple hops through the underlying network

You can use this feature to test the throughput of an Ethernet port remotely. You can verify the maximum rate of frame transmission with no frame loss.

This feature allows for bidirectional or unidirectional throughput measurement, and on-demand or out-of-service (intrusive) operation during service turn-up.

Two types of Ethernet loopback are supported:

- External loopback - Traffic loopback occurs at the Ingress interface. Traffic does not flow into the router for loopback.
- Internal loopback - Traffic loopback occurs at the Egress interface. Traffic loopback occurs after the traffic flows into the router to the other interface.

Ethernet data traffic can be looped back on per port basis. This feature supports a maximum of 100 concurrent Ethernet data plane loopback sessions per system. Filters based on frame header can be used for initiating the loopback session. This ensures that only a subset of traffic that is received on an interface is looped back. You can use Source MAC, Destination MAC, and VLAN Priority (COS bits) as filters.

- N540-28Z4C-SYS-A
- N540-28Z4C-SYS-D
- N540X-16Z4G8Q2C-A
- N540X-16Z4G8Q2C-D
- N540-12Z20G-SYS-A
- N540-12Z20G-SYS-D
- N540X-12Z16G-SYS-A
- N540X-12Z16G-SYS-D

Ethernet Data Plane Loopback Configuration Restrictions

These configuration restrictions are applicable for Ethernet Data Plane Loopback:

- Ethernet data plane loopback is not supported on L3 interfaces or L3 sub-interfaces.
 - The following filters are not supported:
 - Outer VLAN or range of outer VLAN
 - Inner VLAN or range of inner VLAN
 - Ether type
 - Only the following combinations of filters are supported for external loopback:
 - Source MAC
 - Source MAC and Destination MAC
 - Source MAC, Destination MAC, and VLAN priority
 - Destination MAC
 - Destination MAC and VLAN priority
 - The rewrite modification on the loopback traffic is not supported.
 - Ethernet data plane loopback is not supported on packets with destination address as the broadcast MAC address.
 - Ethernet data plane loopback is not supported on BVI interface.
 - Ethernet data plane loopback is not supported on bridge-domain interfaces in Cisco IOS XR Release 6.3.2.
- Layer2 VPN bridge-domains internal loopback is not supported.

- Only one Ethernet loopback session, either internal or external, can be active on the same interface at any given instance.
- This feature supports a maximum throughput of 10Gbps for internal loopback over all the sessions. For external loopback, there is no throughput limit.
- Dropping of packets that are received in the non-loopback direction is not supported.
- Ethernet data plane loopback is not supported on packets having destination as multicast MAC address.
- External and internal Ethernet data plane loopback is not supported over bridge domain.

Configure Ethernet Data Plane Loopback

This section describes how you can configure Ethernet Data Plane Loopback on physical interface and sub-interface. Configuring Ethernet Data Plane Loopback involves these steps:

- Configuring Ethernet Data Plane External Loopback
- Starting an Ethernet Data Plane Loopback Session

Configuration Example

```

/* Configuring Ethernet Data Plane External Loopback */

/* On physical interface */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface tenGigE 0/0/0/0 l2transport
RP/0/RSP0/CPU0:router((config-if-l2)# ethernet loopback permit external

/* Starting an Ethernet Data Plane Loopback Session */

RP/0/RSP0/CPU0:router# ethernet loopback start local interface tenGigE 0/0/0/0 external
source mac-address 0000.0000.0001 destination mac-address 0000.0000.0002 cos 5 timeout none

/* On physical sub-interface */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface tenGigE 0/2/0/0/0.1 l2transport
RP/0/RSP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RSP0/CPU0:router((config-if-l2)# ethernet loopback permit external

/* Starting an Ethernet Data Plane Loopback Session */

RP/0/RSP0/CPU0:router# ethernet loopback start local interface tenGigE 0/2/0/0/0.1 external
source mac-address 0000.0000.0001 destination mac-address 0000.0000.0002 cos 5 timeout
none

/* Configuring Ethernet Data Plane Internal Loopback */

/* On physical interface

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface tenGigE 0/0/0/1 l2transport
RP/0/RSP0/CPU0:router((config-if-l2)# ethernet loopback permit internal

/* Starting an Ethernet Data Plane Loopback Session */

```



```

RP/0/RSP0/CPU0:router# ethernet loopback start local interface tenGigE 0/0/0/1 internal
source mac-address 0000.0000.0002 destination mac-address 0000.0000.0003 cos 5 timeout none

/* On physical sub-interface */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface tenGigE 0/2/0/0/0.1 l2transport
RP/0/RSP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RSP0/CPU0:router(config-if-l2)# ethernet loopback permit internal

/* Starting an Ethernet Data Plane Loopback Session */

RP/0/RSP0/CPU0:router# ethernet loopback start local interface tenGigE 0/2/0/0/0.1 internal
source mac-address 0000.0000.0002 destination mac-address 0000.0000.0003 cos 5 timeout
none

/* Stopping an Ethernet Data Plane Loopback Session */

RP/0/RSP0/CPU0:router# ethernet loopback stop local interface tenGigE 0/0/0/0 id 1
RP/0/RSP0/CPU0:router# ethernet loopback stop local interface tenGigE 0/0/0/1 id 2
RP/0/RSP0/CPU0:router# ethernet loopback stop local interface tenGigE 0/2/0/0/0.1 id 1

```

Similarly, you can configure the Ethernet Data Plane Loopback session for bundle interface and bundle sub-interface.

Running Configuration

This section shows Ethernet Data Plane Loopback running configuration.

```

/* External Loopback */

/* On physical interface */

configure
interface interface tenGigE 0/0/0/0 l2transport
ethernet loopback permit external
!

/* On physical sub-interface */

configure
interface interface tenGigE 0/2/0/0/0.1 l2transport
encapsulation dot1q 100
ethernet loopback permit external
!

/* Internal Loopback */

/* On physical interface */

configure
interface interface tenGigE 0/0/0/1 l2transport
ethernet loopback permit internal
!

```

```

/* On physical sub-interface */

configure
interface interface tenGigE 0/2/0/0/0.1 l2transport
 encapsulation dot1q 100
 ethernet loopback permit internal
!

```

Verification

The following example displays the loopback capabilities per interface. The output shows internal loopback has been permitted on Ten Gigabit Ethernet 0/0/0/1 interface and external loopback has been permitted on Ten Gigabit Ethernet 0/0/0/0 interface.

```
RP/0/RSP0/CPU0:router# show ethernet loopback permitted
```

```

-----
Interface                               Dot1q(s)                               Direction
-----
tenGigE 0/0/0/1.1                       100                                     Internal
tenGigE 0/0/0/0.1                       100                                     External
-----

```

```
/* This example shows all active sessions on the router */
```

```
RP/0/RSP0/CPU0:router# show ethernet loopback active
```

```

Thu Jul 20 11:00:57.864 UTC
Local: TenGigE0/0/0/0.1, ID 1
=====
Direction:                               External
Time out:                                 None
Time left:                                 -
Status:                                    Active
Filters:
  Dot1Q:                                   Any
  Second-dot1Q:                            Any
  Source MAC Address:                      Any
  Destination MAC Address:                 Any
  Class of Service:                        Any
Local: TenGigE0/0/0/0.1, ID 2
=====
Direction:                               External
Time out:                                 None
Time left:                                 -
Status:                                    Active
Filters:
  Dot1Q:                                   Any
  Second-dot1Q:                            Any
  Source MAC Address:                      0000.0000.0001
  Destination MAC Address:                 0000.0000.0002
  Class of Service:                        5

```

Related Topics

- [Ethernet Data Plane Loopback, on page 4](#)

Associated Commands

- ethernet loopback
- show ethernet loopback



CHAPTER 2

L2CP Tunneling MEF

This chapter introduces you to L2 Control Protocols (L2CP) tunneling to help initiate control packets from a local (customer-edge) CE device to a remote CE device.

- [L2CP Tunneling, on page 11](#)
- [Configure L2CP Tunneling, on page 12](#)

L2CP Tunneling

The system supports the following tunnel protocols:


- Link Layer Discovery Protocol (LLDP)
- Link Aggregation Control Protocol (LACP)
- Operation, Administration, Management (OAM)
- Ethernet Local Management Interface (ELMI)
- Cisco Discovery Protocol (CDP)

On a subinterface, when control packets such as LLDP and LACP are tunneled, the system tunnels the same control packets to the main interface.

The LACP packet for VPLS also known as E-LAN service either gets peered or dropped. LACP tunneling is not supported for VPLS service. Tunneling of LACP packets is supported only for VPWS and EVPN-VPWS services.

The router allows to tunnel layer 2 packets between CEs. The following figure depicts Layer 2 Protocol Tunneling. The layer 2 traffic is sent through the S-network, and the S-network switches the traffic from end-to-end. Third-party PE forwards S-tagged frames and peers untagged frames.

Figure 1: L2CP Tunneling



Prerequisites for L2CP Tunneling

A Cisco IOS XR Software that supports Layer 2 Control Protocol Tunneling must be installed previously on the router.

Configure L2CP Tunneling

You do not need to configure L2CP tunneling explicitly. L2CP packets are tunneled over Layer 2 tunnel by default.

Protocol	Packet Type	Action
CDP	Untagged	Peer
LACP	Untagged	Peer
LLDP	Untagged	Peer else Tunnelled
STP	Untagged	Peer
VTP	Untagged	Peer
OAM	Untagged	Peer
BPDU	Untagged	Tunnelled
UDLD	Untagged	Peer
CDP	Tagged	Tunnelled
LACP	Tagged	Tunnelled
LLDP	Tagged	Tunnelled
STP	Tagged	Tunnelled
VTP	Tagged	Tunnelled
BPDU	Tagged	Tunnelled
OAM	Tagged	Tunnelled
ELMI	Tagged	Tunnelled
UDLD	Tagged	Peer



CHAPTER 3

Configure Layer 2 Access Control Lists

This chapter introduces you to Layer 2 Access Control Lists and describe how you can configure the Layer 2 access control lists.

- [Layer 2 Access Control Lists, on page 13](#)
- [Prerequisites for Configuring Layer 2 Access Control Lists, on page 13](#)
- [Layer 2 Access Control Lists Feature Highlights, on page 14](#)
- [Purpose of Layer 2 Access Control Lists, on page 14](#)
- [How a Layer 2 Access Control List Works, on page 14](#)
- [Layer 2 Access Control List Process and Rules, on page 14](#)
- [Create Layer 2 Access Control List, on page 15](#)
- [Restrictions for Configuring Layer 2 Access Control Lists, on page 15](#)
- [Configuration, on page 15](#)

Layer 2 Access Control Lists

An Ethernet services access control lists (ACLs) consist of one or more access control entries (ACE) that collectively define the Layer 2 network traffic profile. This profile can then be referenced by Cisco IOS XR software features. Each Ethernet services ACL includes an action element (permit or deny) based on criteria such as source and destination address, Class of Service (CoS), ether-type, or 802.1ad DEI.

Layer 2 ACLs are supported on ingress traffic only. Layer 2 ACLs are not supported on egress traffic.

Layer 2 access control lists are also known as Ethernet services control access lists.

Prerequisites for Configuring Layer 2 Access Control Lists

This prerequisite applies to configuring the access control lists and prefix lists:

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command.

If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Layer 2 Access Control Lists Feature Highlights

Layer 2 access control lists have these feature highlights:

- The ability to clear counters for an access list using a specific sequence number.
- The ability to copy the contents of an existing access list to another access list.
- Allows users to apply sequence numbers to permit or deny statements.
- Layer 2 ACLs can be applied on interfaces, VLAN subinterfaces, bundle-Ethernet interfaces, bundle subinterfaces with L2 transport. Atomic replacement of Layer 2 ACLs is supported on these physical and bundle interfaces.

Purpose of Layer 2 Access Control Lists

Layer 2 access control lists perform packet filtering to control which packets move through the network and where. Such controls help to limit incoming and outgoing network traffic and restrict the access of users and devices to the network at the port level.

How a Layer 2 Access Control List Works

A Layer 2 access control list is a sequential list consisting of permit and deny statements that apply to Layer 2 configurations. The access list has a name by which it is referenced.

An access list can be configured and named, but it is not in effect until the access list is referenced by a command that accepts an access list. Multiple commands can reference the same access list. An access list can control Layer 2 traffic arriving at the router, but not traffic originating at the router and leaving the router.

Layer 2 Access Control List Process and Rules

Use this process and rules when configuring Layer 2 access control list:

- The software tests the source or destination address of each packet being filtered against the conditions in the access list, one condition (permit or deny statement) at a time.
- If a packet does not match an access list statement, the packet is then tested against the next statement in the list.
- If a packet and an access list statement match, the remaining statements in the list are skipped and the packet is permitted or denied as specified in the matched statement. The first entry that the packet matches determines whether the software permits or denies the packet. That is, after the first match, no subsequent entries are considered.
- If the access list denies the address or protocol, the software discards the packet.
- If no conditions match, the software drops the packet because each access list ends with an unwritten or implicit deny statement. That is, if the packet has not been permitted or denied by the time it was tested against each statement, it is denied.

- The access list should contain at least one permit statement or else all packets are denied.
- Because the software stops testing conditions after the first match, the order of the conditions is critical. The same permit or deny statements specified in a different order could result in a packet being passed under one circumstance and denied in another circumstance.
- Inbound access lists process packets arriving at the router. An inbound access list is efficient because it saves the overhead of routing lookups if the packet is to be discarded because it is denied by the filtering tests. If the packet is permitted by the tests, it is then processed for routing. For inbound lists, permit means continue to process the packet after receiving it on an inbound interface; deny means discard the packet.
- An access list can not be removed if that access list is being applied by an access group in use. To remove an access list, remove the access group that is referencing the access list and then remove the access list.
- An access list must exist before you can use the **ethernet-services access-group** command.

Create Layer 2 Access Control List

Consider these when creating a Layer 2 access control list:

- Create the access list before applying it to an interface.
- Organize your access list so that more specific references appear before more general ones.

Restrictions for Configuring Layer 2 Access Control Lists

These restrictions apply to configuring Layer 2 access control lists:

- Layer 2 access control lists are not supported over management interfaces.
- NetIO (software slow path) is not supported for Layer 2 access control lists.
- Layer 2 access control lists attachment is possible only in ingress direction on an interface.
- Layer 2 access control lists are supported only for the field's L2 source and destination address, Ether Type, Outer VLAN ID, Class of Service (COS), and VLAN Discard Eligibility Indication (DEI). VLAN range is not supported.

Configuration

This section describes how you can configure Layer 2 access control lists.

```
Router# configure
Router(config)# ethernet-services access-list es_acl_1
Router(config-es-acl)# deny 00ff.eedd.0010 ff00.0000.00ff 0000.0100.0001 0000.0000.ffff
Router(config-es-acl)# permit host 000a.000b.000c host 00aa.ab99.1122 cos 1 dei
Router(config-es-acl)# deny host 000a.000b.000c host 00aa.dc11.ba99 cos 7 dei
Router(config-es-acl)# commit
Router(config)# interface tengige0/0/0/4
```

```
Router(config-if)# l2transport
Router(config-if-l2)# commit
Router(config-if-l2)# exit
Router(config-if)# ethernet-services access-group es_acl_1 ingress
Router(config-if)# commit
```

Running Configuration

```
!
Configure
ethernet-services access-list es_acl_1
10 deny 00ff.eedd.0000 ff00.0000.00ff 0000.0100.0000 0000.0000.ffff
20 permit host 000a.000b.000c host 00aa.ab99.1122 cos 1 dei
30 deny host 000a.000b.000c host 00aa.dc11.ba99 cos 7 dei
!
```

Verification

Verify that you have configured Layer 2 access control lists.

```
/* Verify the Layer 2 access control lists configuration */
Router# show access-lists ethernet-services es_acl_1 hardware ingress location 0/0/CPU0
Fri Oct 21 09:39:52.904 UTC
ethernet-services access-list es_acl_1
10 deny 00ff.eedd.0000 ff00.0000.00ff 0000.0100.0000 0000.0000.ffff (2051 matches)
20 permit host 000a.000b.000c host 00aa.ab99.1122 cos 1 dei
30 deny host 000a.000b.000c host 00aa.dc11.ba99 cos 7 dei (2050 matches)
```



CHAPTER 4

Configure Link Bundles for Layer 2 VPNs

An ethernet link bundle is a group of one or more ports that are aggregated together and treated as a single link. Each bundle has a single MAC, a single IP address, and a single configuration set (such as ACLs or QoS).

The advantages of link bundling are:

- Redundancy - Because bundles have multiple links, the failure of a single link does not cause a loss of connectivity.
- Increased bandwidth - On bundled interfaces traffic is forwarded over all available members of the bundle aggregating individual port capacity.

There are two types of link bundling supported depending on the type of interface forming the bundle:

- Ethernet interfaces
- VLAN interfaces (bundle sub-interfaces)

This section describes the configuration of ethernet and VLAN link bundles for use in Layer 2 VPNs.

- [Configure Gigabit Ethernet Link Bundle, on page 17](#)
- [Configure VLAN Bundle, on page 20](#)
- [References for Configuring Link Bundles, on page 21](#)

Configure Gigabit Ethernet Link Bundle

Cisco IOS XR software supports the EtherChannel method of forming bundles of Ethernet interfaces. EtherChannel is a Cisco proprietary technology that allows the user to configure links to join a bundle, but has no mechanisms to check whether the links in a bundle are compatible.

IEEE 802.3ad encapsulation employs a Link Aggregation Control Protocol (LACP) to ensure that all the member links in an ethernet bundle are compatible. Links that are incompatible or have failed are automatically removed from the bundle.

Cisco NCS 540 Series Router supports 100G link bundles.

Restrictions

- All links within a single ethernet link bundle must be configured either to run 802.3ad (LACP) or Etherchannel (non-LACP). Mixed links within a single bundle are not supported.
- MAC accounting is not supported on Ethernet link bundles.

- The maximum number of supported links in each ethernet link bundle is 64.
- The maximum number of supported ethernet link bundles is 128.
- You observe a traffic drop for a few seconds for Layer 2, Layer 3, and BUM traffic when you add bundle members to the existing bundles on the NCS57 line card.


Configuration Example

To create a link bundle between two routers, you must complete the following configurations:

1. Create a bundle instance
2. Map physical interface (s) to the bundle.

Sample values are provided in the following figure.

Figure 2: Link Bundle Topology



For an Ethernet bundle to be active, you must perform the same configuration on both connection endpoints of the bundle.

Configuration

```

/* Enter the global configuration mode and create the ethernet link bundle */
Router# configure
Router(config)# interface Bundle-Ether 3
Router(config-if)# ipv4 address 10.1.2.3 255.0.0.0
Router(config-if)# bundle maximum-active links 32 hot-standby
Router(config-if)# bundle minimum-active links 1
Router(config-if)# bundle minimum-active bandwidth 30000000
Router(config-if)# exit

/* Map physical interfaces to the bundle */
/* Note: Mixed link bundle mode is supported only when active-standby operation is configured
*/
Router(config)# interface TenGigE 1/0/0/0
Router(config-if)# bundle id 3 mode on
Router(config-if)# no shutdown
Router(config)# exit

Router(config)# interface TenGigE 1/0/0/1
Router(config-if)# bundle id 3 mode on
Router(config-if)# no shutdown
Router(config-if)# exit

Router(config)# interface TenGigE 1/0/0/2
Router(config-if)# bundle id 3 mode on
Router(config-if)# no shutdown
Router(config-if)# exit

```

Running Configuration

```
Router# show running-configuration
configure
interface Bundle-Ether 3
  ipv4 address 10.1.2.3 255.0.0.0
  bundle maximum-active links 32 hot-standby
  bundle minimum-active links 1
  bundle minimum-active bandwidth 30000000
!
interface TenGigE 1/0/0/0
  bundle-id 3 mode on
!

interface TenGigE 1/0/0/1
  bundle-id 3 mode on
!

interface TenGigE 1/0/0/2
  bundle-id 3 mode on
!
```

Verification

Verify that interfaces forming the bundle are active and the status of the bundle is Up.

```
Router# show bundle bundle-ether 3
Tue Feb  4 18:24:25.313 UTC
```

```
Bundle-Ether1
```

```
Status: Up
Local links <active/standby/configured>: 3 / 0 / 3
Local bandwidth <effective/available>: 30000000 (30000000) kbps
MAC address (source): 1234.1234.1234 (Configured)
Inter-chassis link: No
Minimum active links / bandwidth: 1 / 1 kbps
Maximum active links: 32
Wait while timer: 2000 ms
Load balancing: Default
LACP: Not operational
  Flap suppression timer: Off
  Cisco extensions: Disabled
  Non-revertive: Disabled
mLACP: Not configured
IPv4 BFD: Not configured
```

Port	Device	State	Port ID	B/W, kbps
Tel1/0/0/0	Local	Active	0x8000, 0x0000	10000000
Link is Active				
Tel1/0/0/1	Local	Active	0x8000, 0x0000	10000000
Link is Active				
Tel1/0/0/2	Local	Active	0x8000, 0x0000	10000000
Link is Active				

Associated Commands

- [bundle maximum-active links](#)
- [interface Bundle-Ether](#)

- [show bundle Bundle-Ether](#)

Configure VLAN Bundle

The procedure for creating VLAN bundle is the same as the procedure for creating VLAN sub-interfaces on a physical ethernet interface.

Configuration Example

To configure VLAN bundles, complete the following configurations:

- Create a bundle instance.
- Create a VLAN interface (bundle sub-interface).
- Map the physical interface(s) to the bundle.

For a VLAN bundle to be active, you must perform the same configuration on both end points of the VLAN bundle.

Configuration

```
/* Enter global configuration mode and create VLAN bundle */
Router# configure
Router(config)# interface Bundle-Ether 2
Router(config-if)# ipv4 address 50.0.0.1/24
Router(config-if)# bundle maximum-active links 32 hot-standby
Router(config-if)# bundle minimum-active bandwidth 30000000
Router(config-if)# bundle minimum-active links 1
Router(config-if)# commit

/* Create VLAN sub-interface and add to the bundle */
Router(config)# interface Bundle-Ether 2.201
Router(config-subif)# ipv4 address 12.22.1.1 255.255.255.0
Router(config-subif)# encapsulation dot1q 201
Router(config-subif)# commit

/* Map the physical interface to the bundle */
Router(config)# interface TenGigE 0/0/0/14
Router(config-if)# bundle id 2 mode on
Router(config-if)# no shutdown
Router(config-if)# commit

/* Repeat the above steps for all the member interfaces:
0/0/0/15, 0/0/0/16 and 0/0/0/17 in this example */
```

Running Configuration

```
configure
interface Bundle-Ether2
  ipv4 address 50.0.0.1 255.255.255.0
  mac-address 1212.1212.1212
  bundle maximum-active links 32 hot-standby
  bundle minimum-active links 1
  bundle minimum-active bandwidth 30000000
!
```

```
interface Bundle-Ether2.201
  ipv4 address 12.22.1.1 255.255.255.0
  encapsulation dot1q 201
  !
interface TenGigE0/0/0/14
  bundle id 2 mode on
  !
interface TenGigE0/0/0/15
  bundle id 2 mode on
  !
interface TenGigE0/0/0/16
  bundle id 2 mode on
  !
interface TenGigE0/0/0/17
  bundle id 2 mode on
  !
```

Verification

Verify that the VLAN status is UP.

```
Router# show interfaces bundle-ether 2.201
```

```
Wed Feb  5 17:19:53.964 UTC
Bundle-Ether2.201 is up, line protocol is up
  Interface state transitions: 1
  Hardware is VLAN sub-interface(s), address is 28c7.ce01.dc7b
  Internet address is 12.22.1.1/24
  MTU 1518 bytes, BW 20000000 Kbit (Max: 20000000 Kbit)
    reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation 802.1Q Virtual LAN, VLAN Id 201,  loopback not set,
  Last link flapped 07:45:25
  ARP type ARPA, ARP timeout 04:00:00
  Last input 00:00:00, output never
  Last clearing of "show interface" counters never
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    2938 packets input, 311262 bytes, 0 total input drops
  - - -
  - - -
```

Associated Commands

- [bundle maximum-active links](#)
- [interface Bundle-Ether](#)
- [show bundle Bundle-Ether](#)

References for Configuring Link Bundles

Characteristics of Link Bundles

- Any type of Ethernet interfaces can be bundled, with or without the use of LACP (Link Aggregation Control Protocol).

- Physical layer and link layer configuration are performed on individual member links of a bundle.
- Configuration of network layer protocols and higher layer applications is performed on the bundle itself.
- A bundle can be administratively enabled or disabled.
- Each individual link within a bundle can be administratively enabled or disabled.
- Ethernet link bundles are created in the same way as Etherchannel channels, where the user enters the same configuration on both end systems.
- The MAC address that is set on the bundle becomes the MAC address of the links within that bundle.
- When LACP configured, each link within a bundle can be configured to allow different keepalive periods on different members.
- Load balancing is done by flow instead of by packet. Data is distributed to a link in proportion to the bandwidth of the link in relation to its bundle.
- QoS is supported and is applied proportionally on each bundle member.
- Link layer protocols, such as CDP, work independently on each link within a bundle.
- Upper layer protocols, such as routing updates and hello messages, are sent over any member link of an interface bundle.
- Bundled interfaces are point to point.
- A link must be in the UP state before it can be in distributing state in a bundle.
- Access Control List (ACL) configuration on link bundles is identical to ACL configuration on regular interfaces.
- Multicast traffic is load balanced over the members of a bundle. For a given flow, internal processes select the member link and all traffic for that flow is sent over that member.

Methods of Forming Bundles of Ethernet Interfaces

Cisco IOS-XR software supports the following methods of forming bundles of Ethernet interfaces:

- IEEE 802.3ad—Standard technology that employs a Link Aggregation Control Protocol (LACP) to ensure that all the member links in a bundle are compatible. Links that are incompatible or have failed are automatically removed from a bundle.

For each link configured as bundle member, information is exchanged between the systems that host each end of the link bundle:

- A globally unique local system identifier
- An identifier (operational key) for the bundle of which the link is a member
- An identifier (port ID) for the link
- The current aggregation status of the link

This information is used to form the link aggregation group identifier (LAG ID). Links that share a common LAG ID can be aggregated. Individual links have unique LAG IDs.

The system identifier distinguishes one router from another, and its uniqueness is guaranteed through the use of a MAC address from the system. The bundle and link identifiers have significance only to the router assigning them, which must guarantee that no two links have the same identifier, and that no two bundles have the same identifier.

The information from the peer system is combined with the information from the local system to determine the compatibility of the links configured to be members of a bundle.

Bundle MAC addresses in the routers come from a set of reserved MAC addresses in the backplane. This MAC address stays with the bundle as long as the bundle interface exists. The bundle uses this MAC address until the user configures a different MAC address. The bundle MAC address is used by all member links when passing bundle traffic. Any unicast or multicast addresses set on the bundle are also set on all the member links.



Note It is recommended that you avoid modifying the MAC address, because changes in the MAC address can affect packet forwarding.

- EtherChannel—Cisco proprietary technology that allows the user to configure links to join a bundle, but has no mechanisms to check whether the links in a bundle are compatible.

Link Aggregation Through LACP

The optional Link Aggregation Control Protocol (LACP) is defined in the IEEE 802 standard. LACP communicates between two directly connected systems (or peers) to verify the compatibility of bundle members. For a router, the peer can be either another router or a switch. LACP monitors the operational state of link bundles to ensure these:

- All links terminate on the same two systems.
- Both systems consider the links to be part of the same bundle.
- All links have the appropriate settings on the peer.

LACP transmits frames containing the local port state and the local view of the partner system's state. These frames are analyzed to ensure both systems are in agreement.



CHAPTER 5

Configure Point-to-Point Layer 2 Services

This section introduces you to point-to-point Layer 2 services, and also describes the configuration procedures to implement it.

The following point-to-point services are supported:

- **Local Switching**—A point-to-point internal circuit on a router, also known as local connect.
- **Attachment circuit**—A connection between a PE-CE router pair.
- **Pseudowires**—A virtual point-to-point circuit from one PE router to another. Pseudowires are implemented over the MPLS network.



Note Point-to-point Layer 2 services are also called as MPLS Layer 2 VPNs.

- [Ethernet over MPLS](#), on page 25
- [Configure Local Switching Between Attachment Circuits](#), on page 28
- [Configure Static Point-to-Point Connections Using Cross-Connect Circuits](#), on page 32
- [Flexible Cross-Connect Service](#), on page 34
- [Flexible Cross-Connect Service Supported Modes](#), on page 35
- [Configure Preferred Tunnel Path](#), on page 49
- [Configure Pseudowire Redundancy](#), on page 50
- [Access Pseudowire Redundancy](#), on page 51
- [Virtual Circuit Connection Verification on L2VPN](#), on page 53
- [GTP Load Balancing](#), on page 53

Ethernet over MPLS

Ethernet-over-MPLS (EoMPLS) provides a tunneling mechanism for Ethernet traffic through an MPLS-enabled Layer 3 core, and encapsulates Ethernet protocol data units (PDUs) inside MPLS packets (using label stacking) to forward them across the MPLS network.

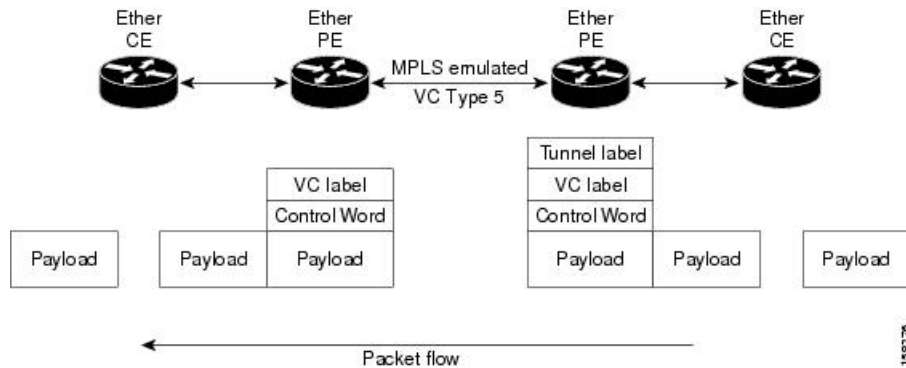
The following sections describe the different modes of implementing EoMPLS.

Ethernet Port Mode

In Ethernet port mode, both ends of a pseudowire are connected to Ethernet ports. In this mode, the port is tunneled over the pseudowire or, using local switching (also known as an *attachment circuit-to-attachment circuit cross-connect*) switches packets or frames from one attachment circuit (AC) to another AC attached to the same PE node.

This figure shows a sample ethernet port mode packet flow:

Figure 3: Ethernet Port Mode Packet Flow

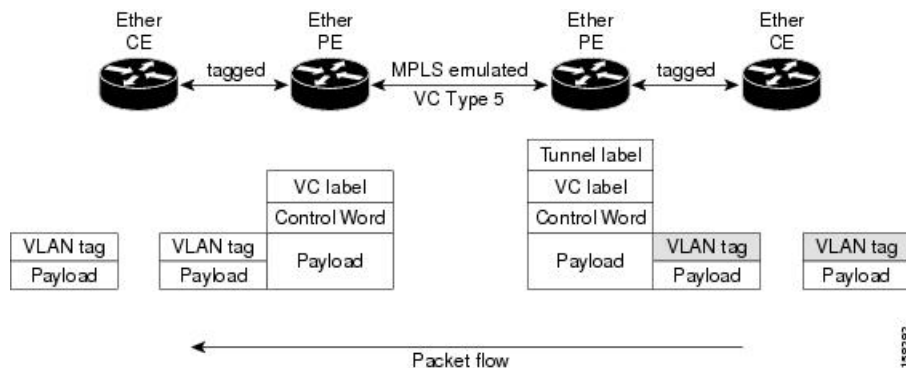


VLAN Mode

In VLAN mode, each VLAN on a customer-end to provider-end link can be configured as a separate L2VPN connection using virtual connection (VC) type 4 or VC type 5. VC type 5 is the default mode.

As illustrated in the following figure, the Ethernet PE associates an internal VLAN-tag to the Ethernet port for switching the traffic internally from the ingress port to the pseudowire; however, before moving traffic into the pseudowire, it removes the internal VLAN tag.

Figure 4: VLAN Mode Packet Flow



At the egress VLAN PE, the PE associates a VLAN tag to the frames coming off of the pseudowire and after switching the traffic internally, it sends out the traffic on an Ethernet trunk port.



Note Because the port is in trunk mode, the VLAN PE doesn't remove the VLAN tag and forwards the frames through the port with the added tag.

QinQ Mode

QinQ is an extension of 802.1Q for specifying multiple 802.1Q tags (IEEE 802.1QinQ VLAN Tag stacking). Layer 3 VPN service termination and L2VPN service transport are enabled over QinQ sub-interfaces.

Cisco NCS 540 Series Routers implement the Layer 2 tunneling or Layer 3 forwarding depending on the sub-interface configuration at provider edge routers. This function only supports up to two QinQ tags on the router:

- Layer 2 QinQ VLANs in L2VPN attachment circuit: QinQ L2VPN attachment circuits are configured under the Layer 2 transport sub-interfaces for point-to-point EoMPLS based cross-connects using both virtual circuit type 4 and type 5 pseudowires and point-to-point local-switching-based cross-connects including full inter-working support of QinQ with 802.1q VLANs and port mode.
- Layer 3 QinQ VLANs: Used as a Layer 3 termination point, both VLANs are removed at the ingress provider edge and added back at the remote provider edge as the frame is forwarded.


Layer 3 services over QinQ include:

- IPv4 unicast and multicast
- IPv6 unicast and multicast
- MPLS
- Connectionless Network Service (CLNS) for use by Intermediate System-to-Intermediate System (IS-IS) Protocol

In QinQ mode, each CE VLAN is carried into an SP VLAN. QinQ mode should use VC type 5, but VC type 4 is also supported. On each Ethernet PE, you must configure both the inner (CE VLAN) and outer (SP VLAN).

The following figure illustrates QinQ using VC type 4.

Figure 5: EoMPLS over QinQ Mode



Note EoMPLS does not support pseudowire stitching or multi segments.

Configure Local Switching Between Attachment Circuits

Local switching involves the exchange of L2 data from one attachment circuit (AC) to the other, and between two interfaces of the same type on the same router. The two ports configured in a local switching connection form an attachment circuit (AC). A local switching connection works like a bridge domain that has only two bridge ports, where traffic enters from one port of the local connection and leaves through the other.

These are some of the characteristics of Layer 2 local switching:

- Layer 2 local switching uses Layer 2 MAC addresses instead of the Layer 3 IP addresses.
- Because there is no bridging involved in a local connection, there is neither MAC learning nor flooding.
- Unlike in a bridge domain, the ACs in a local connection are not in the UP state if the interface state is DOWN.
- Local switching ACs utilize a full variety of Layer 2 interfaces, including Layer 2 trunk (main) interfaces, bundle interfaces, and EFPs.
- Same-port local switching allows you to switch Layer 2 data between two circuits on the same interface.

Restrictions


- All sub-interfaces under the given physical port support only two Tag Protocol Identifiers (TPIDs), such as:
 - 0x88a8, 0x8100
 - 0x9100, 0x8100
 - 0x9200, 0x8100
- VLAN and TPID-based ingress packet filtering is not supported.
- Egress TPID rewrite is not supported.

Topology

An Attachment Circuit (AC) binds a Customer Edge (CE) router to a Provider Edge (PE) router. The PE router uses a pseudowire over the MPLS network to exchange routes with a remote PE router. To establish a point-to-point connection in a Layer 2 VPN from one Customer Edge (CE) router to another (remote router), a mechanism is required to bind the attachment circuit to the pseudowire. A Cross-Connect Circuit (CCC) is used to bind attachment circuits to pseudowires to emulate a point-to-point connection in a Layer 2 VPN.

The following topology is used for configuration.

Figure 6: Local Switching Between Attachment Circuits



Configuration

To configure an AC-AC local switching, complete the following configuration:

- Enable Layer 2 transport on main interfaces.
- Create sub-interfaces with Layer 2 transport enabled, and specify the respective encapsulation for each.
- Enable local switching between the main interfaces, and between the sub-interfaces.
 - Create a cross-connect group.
 - Create a point-to-point cross connect circuit (CCC).
 - Assign interface(s) to the point-to-point cross connect group.

```

/* Enter the interface configuration mode and configure
   L2 transport on the TenGigE interfaces */
Router# configure
Router(config)# interface TenGigE 0/0/0/1 l2transport
Router(config-if-l2)# no shutdown
Router(config-if)# exit
Router(config)# interface TenGigE 0/0/0/9 l2transport
Router(config-if-l2)# no shutdown
Router(config-if-l2)# commit

/* Configure L2 transport and encapsulation on the VLAN sub-interfaces */
Router# configure
Router(config)# interface TenGigE 0/0/0/0.1 l2transport
Router(config-subif)# encapsulation dot1q 5
Router(config-subif)# exit
Router(config)# interface TenGigE 0/0/0/8.1 l2transport
Router(config-subif)# encapsulation dot1q 5
Router(config-subif)# commit

/* Configure ethernet link bundles */
Router# configure
Router(config)# interface Bundle-Ether 3
Router(config-if)# ipv4 address 10.1.3.3 255.0.0.0
Router(config-if)# bundle maximum-active links 32 hot-standby
Router(config-if)# bundle minimum-active links 1
Router(config-if)# bundle minimum-active bandwidth 30000000
Router(config-if)# exit

Router(config)# interface Bundle-Ether 2
Router(config-if)# ipv4 address 10.1.2.2 255.0.0.0
Router(config-if)# bundle maximum-active links 32 hot-standby
Router(config-if)# bundle minimum-active links 1
Router(config-if)# bundle minimum-active bandwidth 30000000
Router(config-if)# exit

/* Add physical interfaces to the ethernet link bundles */
Router(config)# interface TenGigE 0/0/0/1
Router(config-if)# bundle id 3 mode on
Router(config-if)# no shutdown
Router(config)# exit
Router(config)# interface TenGigE 0/0/0/2
Router(config-if)# bundle id 3 mode on
Router(config-if)# no shutdown
Router(config)# exit

```

```

Router(config)# interface TenGigE 0/0/0/9
Router(config-if)# bundle id 2 mode on
Router(config-if)# no shutdown
Router(config-if)# exit
Router(config)# interface TenGigE 0/0/0/8
Router(config-if)# bundle id 2 mode on
Router(config-if)# no shutdown
Router(config-if)# exit

/* Configure Layer 2 transport on the ethernet link bundles */
Router(config)# interface Bundle-Ether 3 l2transport
Router(config-if-l2)# no shutdown
Router(config-if)# exit
Router(config)# interface Bundle-Ether 2 l2transport
Router(config-if-l2)# no shutdown
Router(config-if-l2)# commit

/* Configure local switching on the TenGigE Interfaces */
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group XCON1
Router(config-l2vpn-xc)# p2p XCON1_P2P3
Router(config-l2vpn-xc-p2p)# interface TenGigE0/0/0/1
Router(config-l2vpn-xc-p2p)# interface TenGigE0/0/0/9
Router(config-l2vpn-xc-p2p)# commit
Router(config-l2vpn-xc-p2p)# exit

/* Configure local switching on the VLAN sub-interfaces */
Router(config-l2vpn-xc)# p2p XCON1_P2P1
Router(config-l2vpn-xc-p2p)# interface TenGigE0/0/0/0.1
Router(config-l2vpn-xc-p2p)# interface TenGigE0/0/0/8.1
Router(config-l2vpn-xc-p2p)# commit
Router(config-l2vpn-xc-p2p)# exit

/* Configure local switching on ethernet link bundles */
Router(config-l2vpn-xc)# p2p XCON1_P2P4
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether 3
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether 2
Router(config-l2vpn-xc-p2p)# commit

```

Running Configuration

```

configure
 interface tenGigE 0/0/0/1 l2transport
 !
 interface tenGigE 0/0/0/9 l2transport
 !
 !

 interface tenGigE 0/0/0/0.1 l2transport
 encapsulation dot1q 5
 rewrite ingress tag push dot1q 20 symmetric
 !
 interface tenGigE 0/0/0/8.1 l2transport
 encapsulation dot1q 5
 !
 interface Bundle-Ether 3 l2transport
 !
 interface Bundle-Ether 2 l2transport
 !

```



```

l2vpn
xconnect group XCON1
  p2p XCON1_P2P3
    interface TenGigE0/0/0/1
    interface TenGigE0/0/0/9
    !
  !
!
l2vpn
xconnect group XCON1
  p2p XCON1_P2P1
    interface TenGigE0/0/0/0.1
    interface TenGigE0/0/0/8.1
    !
  !
!
l2vpn
xconnect group XCON1
  p2p XCON1_P2P4
    interface Bundle-Ether 3
    interface Bundle-Ether 2
    !
  !
!

```

Verification

- Verify if the configured cross-connect is UP

```
router# show l2vpn xconnect brief
```

Locally Switching

Like-to-Like	UP	DOWN	UNR
EFP	1	0	0
Total	1	0	0
Total	1	0	0

Total: 1 UP, 0 DOWN, 0 UNRESOLVED

```
router# show l2vpn xconnect
```

Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

XConnect Group	Name	ST	Segment 1 Description	ST	Segment 2 Description	ST
XCON1	XCON_P2P1	UP	Te0/0/0/1	UP	Te0/0/0/9	UP
XCON1	XCON_P2P3	UP	Te0/0/0/0.1	UP	Te0/0/0/8.1	UP

Associated Commands

- `interface (p2p)`
- `l2vpn`
- `p2p`
- `xconnect group`

Configure Static Point-to-Point Connections Using Cross-Connect Circuits

This section describes how you can configure static point-to-point cross connects in a Layer 2 VPN.

Requirements and Limitations

Before you can configure a cross-connect circuit in a Layer 2 VPN, ensure that the following requirements are met:

- The CE and PE routers are configured to operate in the MPLS network.
- The name of a cross-connect circuit is configured to identify a pair of PE routers and must be unique within the cross-connect group.
- A segment (an attachment circuit or pseudowire) is unique and can belong only to a single cross-connect circuit.
- A static virtual circuit local label is globally unique and can be used in only one pseudowire.
- A maximum of 4000 cross-connects can be configured per PE router.




Note Static pseudowire connections do not use LDP for signaling.

Topology

The following topology is used to configure static cross-connect circuits in a Layer 2 VPN.

Figure 7: Static Cross-Connect Circuits in a Layer 2 VPN



Configuration

```

/* Configure PE1 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group XCON1
Router(config-l2vpn-xc)# p2p xc1
Router(config-l2vpn-xc-p2p)# interface gigabitethernet0/1/0/0.1
Router(config-l2vpn-xc-p2p)# neighbor 10.165.100.151 pw-id 100
Router(config-l2vpn-xc-p2p-pw)# mpls static label local 50 remote 40
Router(config-l2vpn-xc-p2p-pw)# commit

/*Configure PE2 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group XCON1
Router(config-l2vpn-xc)# p2p xc1
Router(config-l2vpn-xc-p2p)# interface gigabitethernet0/2/0/0.4
Router(config-l2vpn-xc-p2p)# neighbor 10.165.200.254 pw-id 100
Router(config-l2vpn-xc-p2p-pw)# mpls static label local 40 remote 50
Router(config-l2vpn-xc-p2p-pw)# commit

```

Running Configuration

```

/* On PE1 */
!
l2vpn
xconnect group XCON1
  p2p xc1
    interface GigabitEthernet0/1/0/0.1
      neighbor ipv4 10.165.100.151 pw-id 100
      mpls static label local 50 remote 40
!

/* On PE2 */
!
l2vpn
xconnect group XCON2
  p2p xc1
    interface GigabitEthernet0/2/0/0.4
      neighbor ipv4 10.165.200.254 pw-id 100
      mpls static label local 40 remote 50
!

```

Verification

```

/* Verify the static cross connect on PE1 */
Router# show l2vpn xconnect
Tue Apr 12 20:18:02.971 IST
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
        SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

```

XConnect		Segment 1		Segment 2		
Group	Name	ST	Description	ST	Description	ST
XCON1	xc1	UP	Gi0/1/0/0.1	UP	10.165.100.151 100	UP

```

/* Verify the static cross connect on PE2 */

```

```

Router# show l2vpn xconnect
Tue Apr 12 20:18:02.971 IST
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,

```

SB = Standby, SR = Standby Ready, (PP) = Partially Programmed

XConnect Group	Name	ST	Segment 1		Segment 2		
			Description	ST	Description	ST	
XCON2	xc1	UP	Gi0/2/0/0.4	UP	10.165.200.254	100	UP


Flexible Cross-Connect Service

The flexible cross-connect service feature enables aggregation of attachment circuits (ACs) across multiple endpoints in a single Ethernet VPN Virtual Private Wire Service (EVPN-VPWS) service instance, on the same Provider Edge (PE). ACs are represented either by a single VLAN tag or double VLAN tags. The associated AC with the same VLAN tag(s) on the remote PE is cross-connected. The VLAN tags define the matching criteria to be used in order to map the frames on an interface to the appropriate service instance. As a result, the VLAN rewrite value must be unique within the flexible cross-connect (FXC) instance to create the lookup table. The VLAN tags can be made unique using the rewrite configuration. The lookup table helps determine the path to be taken to forward the traffic to the corresponding destination AC. This feature reduces the number of tunnels by muxing VLANs across many interfaces. It also reduces the number of MPLS labels used by a router. This feature supports both single-homing and multi-homing.

Flexible Cross-Connect Service - Single-Homed

Consider the following topology in which the traffic flows from CE1 and CE2 to PE1 through ACs. ACs are aggregated across multiple endpoints on the same PE. The VLAN (rewrite) creates the lookup table based on the rewrite configured at AC interfaces on PE1. PE1 uses BGP to exchange routes with PE2 and creates a tunnel over EVPN MPLS network. The VLANs (rewrite) on PE2 must match the rewrite configured on PE1. Based on the rewrite tag, the PE2 forwards the traffic to the corresponding ACs. For example, if the ACs for CE1 and CE3 are configured with the same rewrite tag, the end-to-end traffic is sent from CE1 to CE3.

Figure 8: Flexible Cross-Connect Service




Flexible Cross-Connect Service - Multi-Homed

The Flexible Cross-Connect Service multihoming capability enables you to connect a customer edge (CE) device to two or more provider edge (PE) devices to provide load balancing and redundant connectivity. Flow-based load balancing is used to send the traffic between PEs and CEs. Flow-based load balancing is used to connect source and remote PEs as well. The customer edge device is connected to PE through Ethernet bundle interface.

When a CE device is multi-homed to two or more PEs and when all PEs can forward traffic to and from the multi-homed device for the VLAN, then such multihoming is referred to as all-active multihoming.

Figure 9: Flexible Cross-Connect Service Multi-Homed



Consider the topology in which CE1 and CE2 are multi-homed to PE1 and PE2; CE3 and CE4 are multi-homed to PE3 and PE4. PE1 and PE2 advertise Ethernet A-D Ethernet Segment (ES-EAD) route to remote PEs that is PE3 and PE4. Similarly, PE3 and PE4 advertise ES-EAD route to remote PEs that is PE1 and PE2. The ES-EAD route is advertised per main interface.

Consider a traffic flow from CE1 to CE3. Traffic is sent to either PE1 or PE2. The selection of path is dependent on the CE implementation for forwarding over a LAG. Traffic is encapsulated at each PE and forwarded to the remote PEs (PE 3 and PE4) through the MPLS tunnel. Selection of the destination PE is established by flow-based load balancing. PE3 and PE4 send the traffic to CE3. The selection of path from PE3 or PE4 to CE3 is established by flow-based load balancing.

Flexible Cross-Connect Service Supported Modes

The Flexible Cross-Connect Service feature supports the following modes:

- VLAN Unaware
- VLAN Aware
- Local Switching

VLAN Unaware

In this mode of operation, a group of normalized ACs on a single ES that are destined to a single endpoint or interface are multiplexed into a single EVPN VPWS tunnel represented by a single VPWS service ID. The VLAN-Unaware FXC reduces the number of BGP states. VLAN failure is not signaled over BGP. One EVI/EAD route is advertised per VLAN-Unaware FXC rather than per AC. In multihoming scenario, there will be ES-EAD route as well. EVI can be shared with other VLAN-Unaware FXC or EVPN VPWS. If AC goes down on PE1, the remote PE is not be informed of the failure, and PE3 or PE4 continues to send the traffic to PE1 and PE2 resulting in packet drop.

Multihoming is supported on VLAN Unaware FXC only if all ACs belong to the same main interface.

If you have multiple ESIs, regardless of whether it is a zero-ESI or non-zero ESI, only ESI 0 is signalled. Only single-home mode is supported in this scenario.

Configure Single-Homed Flexible Cross-Connect Service using VLAN Unaware

This section describes how you can configure single-homed flexible cross-connect service using VLAN unaware

```

/* Configure PE1 */
Router# configure
Router(config)# interface GigabitEthernet 0/2/0/3.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100
symetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface GigabitEthernet 0/2/0/0.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200
symetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxs1
Router(config-l2vpn-fxs-vu)# interface GigabitEthernet 0/2/0/3.1
Router(config-l2vpn-fxs-vu)# interface GigabitEthernet 0/2/0/0.1
Router(config-l2vpn-fxs-vu)# neighbor evpn evi 1 target 1
Router(config-l2vpn-fxs-vu)# commit

/* Configure PE2 */
Router# configure
Router(config)# interface GigabitEthernet 0/0/0/3.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100
symetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface GigabitEthernet 0/0/0/0.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200
symetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxs1
Router(config-l2vpn-fxs-vu)# interface GigabitEthernet 0/0/0/3.1
Router(config-l2vpn-fxs-vu)# interface GigabitEthernet 0/0/0/0.1
Router(config-l2vpn-fxs-vu)# neighbor evpn evi 1 target 1
Router(config-l2vpn-fxs-vu)# commit

```

Running Configuration

```

/* On PE1 */
!
Configure
interface GigabitEthernet 0/2/0/3.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100 symetric
!

Configure
interface GigabitEthernet 0/2/0/0.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200 symetric
!

```

```

l2vpn
  flexible-xconnect-service vlan-unaware fxs1
  interface GigabitEthernet 0/2/0/3.1
  interface GigabitEthernet0/2/0/0.1
  neighbor evpn evi 1 target 1

!

/* On PE2 */
!
Configure
interface GigabitEthernet 0/0/0/3.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100 symmetric
!

Configure
interface GigabitEthernet 0/0/0/0.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200 symmetric
!

l2vpn
  flexible-xconnect-service vlan-unaware fxs1
  interface GigabitEthernet 0/0/0/3.1
  interface GigabitEthernet0/0/0/0.1
  neighbor evpn evi 1 target 1

!

```

Configure Multi-Homed Flexible Cross-Connect Service using VLAN Unaware

This section describes how you can configure multi-homed flexible cross-connect service using VLAN unaware.

```

/* Configure PE1 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxc1_16
Router(config-l2vpn-fxs)# interface Bundle-Ether10.11
Router(config-l2vpn-fxs)# interface Bundle-Ether10.12
Router(config-l2vpn-fxs)# neighbor evpn evi 1 target 16
Router(config-l2vpn-fxs)# commit
Router(config-l2vpn-fxs)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether10.11 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface Bundle-Ether10.12 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-subif)# commit
Router(config-subif)# exit
Router(config)# evpn
Router (config-evpn)# interface Bundle-Ether10
Router (config-evpn-ac)# ethernet-segment
Router (config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.0a.00
Router (config-evpn-ac-es)# commit

```

```

/* Configure PE2 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxc1_16
Router(config-l2vpn-fxs-vu)# interface Bundle-Ether10.11
Router(config-l2vpn-fxs)# interface Bundle-Ether10.12
Router(config-l2vpn-fxs)# neighbor evpn evi 1 target 16
Router(config-l2vpn-fxs)# commit
Router(config-l2vpn-fxs)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether10.11 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface Bundle-Ether10.12 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-subif)# commit
Router(config-subif)# exit
Router(config)# evpn
Router (config-evpn)# interface Bundle-Ether10
Router (config-evpn-ac)# ethernet-segment
Router (config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.0a.00
Router (config-evpn-ac-es)# commit

/* Configure PE3 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxc1_16
Router(config-l2vpn-fxs-vu)# interface Bundle-Ether20.11
Router(config-l2vpn-fxs)# interface Bundle-Ether20.12
Router(config-l2vpn-fxs)# neighbor evpn evi 1 target 16
Router(config-l2vpn-fxs)# commit
Router(config-l2vpn-fxs)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether20.11 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif)# commit
Router(config-subif)# exit
Router(config)# interface Bundle-Ether20.12 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-l2vpn-subif)# commit
Router(config-subif)# exit
Router(config)# evpn
Router (config-evpn)# interface Bundle-Ether20
Router (config-evpn-ac)# ethernet-segment
Router (config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.14.00
Router (config-evpn-ac-es)# commit

/* Configure PE4 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxc1_16
Router(config-l2vpn-fxs-vu)# interface Bundle-Ether20.11
Router(config-l2vpn-fxs)# interface Bundle-Ether20.12
Router(config-l2vpn-fxs)# neighbor evpn evi 1 target 16
Router(config-l2vpn-fxs)# commit
Router(config-l2vpn-fxs)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether20.11 l2transport

```



```

Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif)# commit
Router(config-subif)# exit
Router(config)# interface Bundle-Ether20.12 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-l2vpn-subif)# commit
Router(config-subif)# exit
Router(config)# evpn
Router (config-evpn)# interface Bundle-Ether20
Router (config-evpn-ac)# ethernet-segment
Router (config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.14.00
Router (config-evpn-ac-es)# commit

```

Running Configuration

```

/* On PE1 */

configure
l2vpn
flexible-xconnect-service vlan-unaware fxc1_16
interface Bundle-Ether10.11
interface Bundle-Ether10.12
neighbor evpn evi 1 target 16

!

configure
interface Bundle-Ether10.11 l2transport
encapsulation dot1q 1
rewrite ingress tag translate 1-to-1 dot1q 11 symmetric

!

configure
interface Bundle-Ether10.12 l2transport
encapsulation dot1q 2
rewrite ingress tag translate 1-to-1 dot1q 12 symmetric

!

evpn
interface Bundle-Ether10
ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.0a.00

!

/* On PE2 */

configure
l2vpn
flexible-xconnect-service vlan-unaware fxc1_16
interface Bundle-Ether10.11
interface Bundle-Ether10.12
neighbor evpn evi 1 target 16

!

configure
interface Bundle-Ether10.11 l2transport
encapsulation dot1q 1

```

```

rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
!

configure
interface Bundle-Ether10.12 l2transport
 encapsulation dot1q 2
 rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
!

evpn
interface Bundle-Ether10
 ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.0a.00
!

/* On PE3 */

configure
l2vpn
flexible-xconnect-service vlan-unaware fxc1_16
interface Bundle-Ether20.11
interface Bundle-Ether20.12
neighbor evpn evi 1 target 16
!

configure
interface Bundle-Ether20.11 l2transport
 encapsulation dot1q 1
 rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
!

configure
interface Bundle-Ether20.12 l2transport
 encapsulation dot1q 2
 rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
!

evpn
interface Bundle-Ether20
 ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.14.00
!

/* On PE4 */

configure
l2vpn
flexible-xconnect-service vlan-unaware fxc1_16
interface Bundle-Ether20.11
interface Bundle-Ether20.12
neighbor evpn evi 1 target 16
!

configure
interface Bundle-Ether20.11 l2transport
 encapsulation dot1q 1
 rewrite ingress tag translate 1-to-1 dot1q 11 symmetric

```

```

!
configure
interface Bundle-Ether20.12 l2transport
    encapsulation dot1q 2
    rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
!
evpn
interface Bundle-Ether20
    ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.14.00
!

```

VLAN Aware

In this mode of operation, normalized ACs across different Ethernet segments and interfaces are multiplexed into a single EVPN VPWS service tunnel. This single tunnel is represented by many VPWS service IDs (one per normalized VLAN ID (VID)) and these normalized VIDs are signaled using EVPN BGP. The VLAN-Aware FXC reduces the number of PWs; but it does not reduce the BGP states. VLAN failure is signaled over BGP. The VLAN-Aware FXC advertises one EAD route per AC rather than per FXC. For VLAN-Aware FXC, the EVI must be unique to the FXC itself. It cannot be shared with any other service such as FXC, EVPN, EVPN-VPWS, PBB-EVPN. If a single AC goes down on PE1, it withdraws only the EAD routes associated with that AC. The ES-EAD route will also be withdrawn on failure of the main interface. The equal-cost multipath (ECMP) on PE3 or PE4 stops sending traffic for this AC to PE1, and only sends it to PE2.

For the same VLAN-Aware FXC, you can either configure all non-zero ESIs or all zero-ESIs. You cannot configure both zero-ESI and non-zero ESI for the same VLAN-Aware FXC. This applies only to single-home mode.

Configure Single-Homed Flexible Cross-Connect using VLAN Aware

This section describes how you can configure single-homed flexible cross-connect service using VLAN aware.

```

/* Configure PE1 */
Router# configure
Router(config)# interface GigabitEthernet 0/2/0/7.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface GigabitEthernet 0/2/0/7.2 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 4
Router(config-l2vpn-fxs-va)# interface GigabitEthernet 0/2/0/7.1
Router(config-l2vpn-fxs-va)# interface GigabitEthernet 0/2/0/7.2
Router(config-l2vpn-fxs-va)# commit

/* Configure PE2 */
Router# configure
Router(config)# interface GigabitEthernet 0/0/0/7.1 l2transport

```

```

Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100
symetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface GigabitEthernet 0/0/0/7.2 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200
symetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 4
Router(config-l2vpn-fxs-va)# interface GigabitEthernet 0/0/0/7.1
Router(config-l2vpn-fxs-va)# interface GigabitEthernet 0/0/0/7.2
Router(config-l2vpn-fxs-va )# commit

```

Running Configuration

```

/* On PE1 */
!
Configure
interface GigabitEthernet 0/2/0/7.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100 symetric
!

Configure
interface GigabitEthernet 0/2/0/7.2 l2transport
  encapsulation dot1q 2
  rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200 symetric
!

l2vpn
  flexible-xconnect-service vlan-aware evi 4
  interface GigabitEthernet 0/2/0/7.1
  interface GigabitEthernet 0/2/0/7.2

!

/* On PE2 */
!
Configure
interface GigabitEthernet 0/0/0/7.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 500 second-dot1q 100 symetric
!

Configure
interface GigabitEthernet 0/0/0/7.2 l2transport
  encapsulation dot1q 2
  rewrite ingress tag translate 1-to-2 dot1q 600 second-dot1q 200 symetric
!

l2vpn
  flexible-xconnect-service vlan-aware evi 4
  interface GigabitEthernet 0/0/0/7.1
  interface GigabitEthernet 0/0/0/7.2

!

```

Configure Multi-Homed Flexible Cross-Connect Service using VLAN Aware

This section describes how you can configure multi-homed flexible cross-connect service using VLAN aware.

```

/* Configure PE1 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 6
Router(config-l2vpn-fxs-va)# interface Bundle-Ether2.1
Router(config-l2vpn-fxs-va)# interface Bundle-Ether3.1
Router(config-l2vpn-fxs-va)# commit
Router(config-l2vpn-fxs-va)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether2.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface Bundle-Ether3.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether2
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 22.33.44.55.66.77.88.99.aa
Router(config-evpn-ac-es)# commit
Router(config-evpn-ac-es)# exit
Router(config-evpn-ac)# exit
Router(config-evpn)# interface Bundle-Ether3
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 33.44.55.66.77.88.99.aa.bb
Router(config-evpn-ac-es)# commit

/* Configure PE2 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 6
Router(config-l2vpn-fxs-va)# interface Bundle-Ether2.1
Router(config-l2vpn-fxs-va)# interface Bundle-Ether3.1
Router(config-l2vpn-fxs-va)# commit
Router(config-l2vpn-fxs-va)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether2.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface Bundle-Ether3.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 2
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether2
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 22.33.44.55.66.77.88.99.aa
Router(config-evpn-ac-es)# commit
Router(config-evpn-ac-es)# exit
Router(config-evpn-ac)# exit
Router(config-evpn)# interface Bundle-Ether3
Router(config-evpn-ac)# ethernet-segment

```

```

Router(config-evpn-ac-es) # identifier type 0 33.44.55.66.77.88.99.aa.bb
Router(config-evpn-ac-es) # commit

/* Configure PE3 */
Router# configure
Router(config) # l2vpn
Router(config-l2vpn) # flexible-xconnect-service vlan-aware evi 6
Router(config-l2vpn-fxs-va) # interface Bundle-Ether4.1
Router(config-l2vpn-fxs-va) # interface Bundle-Ether5.1
Router(config-l2vpn-fxs-va) # commit
Router(config-l2vpn-fxs-va) # exit
Router(config-l2vpn) # exit
Router(config) # interface Bundle-Ether4.1 l2transport
Router(config-l2vpn-subif) # encapsulation dot1q 1
Router(config-l2vpn-subif) # rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif) # commit
Router(config-l2vpn-subif) # exit
Router(config) # interface Bundle-Ether5.1 l2transport
Router(config-l2vpn-subif) # encapsulation dot1q 2
Router(config-l2vpn-subif) # rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-l2vpn-subif) # commit
Router(config-l2vpn-subif) # exit
Router(config) # evpn
Router(config-evpn) # interface Bundle-Ether4
Router(config-evpn-ac) # ethernet-segment
Router(config-evpn-ac-es) # identifier type 0 00.01.00.ac.ce.55.00.14.00
Router(config-evpn-ac-es) # commit
Router(config-evpn-ac-es) # exit
Router(config-evpn-ac) # exit
Router(config-evpn) # interface Bundle-Ether5
Router(config-evpn-ac) # ethernet-segment
Router(config-evpn-ac-es) # identifier type identifier type 0 00.01.00.ac.ce.55.00.15.00
Router(config-evpn-ac-es) # commit

/* Configure PE4 */
Router# configure
Router(config) # l2vpn
Router(config-l2vpn) # flexible-xconnect-service vlan-aware evi 6
Router(config-l2vpn-fxs-va) # interface Bundle-Ether4.1
Router(config-l2vpn-fxs-va) # interface Bundle-Ether5.1
Router(config-l2vpn-fxs-va) # commit
Router(config-l2vpn-fxs-va) # exit
Router(config-l2vpn) # exit
Router(config) # interface Bundle-Ether4.1 l2transport
Router(config-l2vpn-subif) # encapsulation dot1q 1
Router(config-l2vpn-subif) # rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
Router(config-l2vpn-subif) # commit
Router(config-l2vpn-subif) # exit
Router(config) # interface Bundle-Ether5.1 l2transport
Router(config-l2vpn-subif) # encapsulation dot1q 2
Router(config-l2vpn-subif) # rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
Router(config-l2vpn-subif) # commit
Router(config-l2vpn-subif) # exit
Router(config) # evpn
Router(config-evpn) # interface Bundle-Ether4
Router(config-evpn-ac) # ethernet-segment
Router config-evpn-ac-es) # identifier type 0 00.01.00.ac.ce.55.00.14.00
Router(config-evpn-ac-es) # commit
Router(config-evpn-ac-es) # exit
Router(config-evpn-ac) # exit
Router(config-evpn) # interface Bundle-Ether5
Router(config-evpn-ac) # ethernet-segment

```

```
Router(config-evpn-ac-es)# identifier type identifier type 0 00.01.00.ac.ce.55.00.15.00
Router(config-evpn-ac-es)# commit
```

Running Configuration

```
/* On PE1 */
!
configure
l2vpn
  flexible-xconnect-service vlan-aware evi 6
  interface Bundle-Ether2.1
  interface Bundle-Ether3.1

!

configure
interface Bundle-Ether2.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-1 dot1q 11 symmetric

!

configure
interface Bundle-Ether3.1 l2transport
  encapsulation dot1q 2
  rewrite ingress tag translate 1-to-1 dot1q 12 symmetric

!

evpn
  interface Bundle-Ether2
    ethernet-segment identifier type 0 22.33.44.55.66.77.88.99.aa
  interface Bundle-Ether3
    ethernet-segment identifier type 0 33.44.55.66.77.88.99.aa.bb

!

/* On PE2 */
!
configure
l2vpn
  flexible-xconnect-service vlan-aware evi 6
  interface Bundle-Ether2.1
  interface Bundle-Ether3.1

!

configure
interface Bundle-Ether2.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-1 dot1q 11 symmetric

!

configure
interface Bundle-Ether3.1 l2transport
  encapsulation dot1q 2
  rewrite ingress tag translate 1-to-1 dot1q 12 symmetric

!

evpn
  interface Bundle-Ether2
    ethernet-segment identifier type 0 22.33.44.55.66.77.88.99.aa
  interface Bundle-Ether3
    ethernet-segment identifier type 0 33.44.55.66.77.88.99.aa.bb
```

```

!
/* On PE3 */
!
configure
l2vpn
flexible-xconnect-service vlan-aware evi 6
interface Bundle-Ether4.1
interface Bundle-Ether5.1
!

configure
interface Bundle-Ether4.1 l2transport
encapsulation dot1q 1
rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
!

configure
interface Bundle-Ether5.1 l2transport
encapsulation dot1q 2
rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
!

evpn
interface Bundle-Ether4
ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.14.00
interface Bundle-Ether5
ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.15.00
!

/* On PE4 */
!
configure
l2vpn
flexible-xconnect-service vlan-aware evi 6
interface Bundle-Ether4.1
interface Bundle-Ether5.1
!

configure
interface Bundle-Ether4.1 l2transport
encapsulation dot1q 1
rewrite ingress tag translate 1-to-1 dot1q 11 symmetric
!

configure
interface Bundle-Ether5.1 l2transport
encapsulation dot1q 2
rewrite ingress tag translate 1-to-1 dot1q 12 symmetric
!

evpn
interface Bundle-Ether4
ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.14.00
interface Bundle-Ether5
ethernet-segment identifier type 0 00.01.00.ac.ce.55.00.15.00
!

```


Local Switching


Traffic between the two ACs is locally switched within the PE when two ACs belonging to different Ethernet Segment have the same normalization VLANs. Local switching is supported only on FXC VLAN-aware.

Consider a topology in which CE1 and CE2 have different Ethernet Segment. However, they both have the same normalized VLANs. Hence, when a traffic is sent from CE1 to CE2, PE1 routes the traffic to CE2 using local switching.

If there is a failure and when the link from CE1 to PE1 goes down, PE1 sends the traffic to PE2 through EVPN pseudowire. Then the PE2 sends the traffic to CE2.

CE1 and CE2 must be on different non-zero ESI.

Figure 10: Local Switching



Configure Multi-Homed Flexible Cross-Connect Service using Local Switching

This section describes how you can configure multi-homed flexible cross-connect service using local switching.

```

/* Configure PE1 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 6
Router(config-l2vpn-fxs-va)# interface Bundle-Ether2.1
Router(config-l2vpn-fxs-va)# interface Bundle-Ether3.1
Router(config-l2vpn-fxs-va)# commit
Router(config-l2vpn-fxs)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether2.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3
symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface Bundle-Ether3.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3
symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether2
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 22.33.44.55.66.77.88.99.aa
Router(config-evpn-ac-es)# commit
Router(config-evpn-ac-es)# exit
Router(config-evpn-ac)# exit
Router(config-evpn)# interface Bundle-Ether3
Router(config-evpn-ac)# ethernet-segment

```

Running Configuration

```

Router(config-evpn-ac-es)# identifier type 0 33.44.55.66.77.88.99.aa.bb
Router(config-evpn-ac-es)# commit

/* Configure PE2 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 6
Router(config-l2vpn-fxs-va)# interface Bundle-Ether2.1
Router(config-l2vpn-fxs-va)# interface Bundle-Ether3.1
Router(config-l2vpn-fxs-va)# commit
Router(config-l2vpn-fxs)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether2.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3
symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# interface Bundle-Ether3.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1
Router(config-l2vpn-subif)# rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3
symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether2
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 22.33.44.55.66.77.88.99.aa
Router(config-evpn-ac-es)# commit
Router(config-evpn-ac-es)# exit
Router(config-evpn-ac)# exit
Router(config-evpn)# interface Bundle-Ether3
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 33.44.55.66.77.88.99.aa.bb
Router(config-evpn-ac-es)# commit

```

Running Configuration

```

/* On PE1 */

configure
l2vpn
  flexible-xconnect-service vlan-aware evi 6
  interface Bundle-Ether2.1
  interface Bundle-Ether3.1

!

configure
interface Bundle-Ether2.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3 symmetric

!

configure
interface Bundle-Ether3.1 l2transport
  encapsulation dot1q 1
  rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3 symmetric

!

evpn
  interface Bundle-Ether2

```

```

    ethernet-segment identifier type 0 22.33.44.55.66.77.88.99.aa
interface Bundle-Ether3
    ethernet-segment identifier type 0 33.44.55.66.77.88.99.aa.bb

!

/* On PE2 */

configure
l2vpn
    flexible-xconnect-service vlan-aware evi 6
    interface Bundle-Ether2.1
    interface Bundle-Ether3.1

!

configure
interface Bundle-Ether2.1 l2transport
    encapsulation dot1q 1
    rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3 symmetric

!

configure
interface Bundle-Ether3.1 l2transport
    encapsulation dot1q 1
    rewrite ingress tag translate 1-to-2 dot1q 3 second-dot1q 3 symmetric

!

evpn
    interface Bundle-Ether2
        ethernet-segment identifier type 0 22.33.44.55.66.77.88.99.aa
    interface Bundle-Ether3
        ethernet-segment identifier type 0 33.44.55.66.77.88.99.aa.bb

!

```

Configure Preferred Tunnel Path

Preferred tunnel path functionality lets you map pseudowires to specific traffic-engineering tunnels. Attachment circuits are cross-connected to specific MPLS traffic engineering tunnel interfaces instead of remote PE router IP addresses (reachable using IGP or LDP).

When using a preferred tunnel path, it is assumed that the traffic engineering tunnel that transports the Layer 2 traffic runs between the two PE routers (that is, its head starts at the imposition PE router and its tail terminates on the disposition PE router).

Configuration

```

/* Enter global configuration mode */
Router# configure
Router(config)# l2vpn

/* Configure pseudowire class name */
Router(config-l2vpn)# pw-class path1

/* Configure MPLS encapsulation for the pseudowire */
Router(config-l2vpn-pwc)# encapsulation mpls

```

```

/* Configure preferred path tunnel settings.
If fallback disable configuration is used, and when
the TE/ tunnel is configured,
if the preferred path goes down,
the corresponding pseudowire can also go down. */

Router(config-l2vpn-pwc-encap-mpls) # preferred-path
                                interface tunnel-te 11 fallback disable

/* Commit your configuration */
Router(config-l2vpn-pwc) # exit
Router(config-l2vpn) # commit

```

Running Configuration

```

Router# show running-configuration
!
l2vpn
  pw-class path1
    encapsulation mpls
    preferred-path interface tunnel-te 11 fallback disable
  !
!
!

```

Configure Pseudowire Redundancy

Pseudowire redundancy allows you to configure your network to detect a failure in the network and reroute the Layer 2 service to another endpoint that can continue to provide service. This feature provides the ability to recover from a failure of either the remote provider edge (PE) router or the link between the PE and customer edge (CE) routers.

L2VPNs can provide pseudowire resiliency through their routing protocols. When connectivity between end-to-end PE routers fails, an alternative path to the directed LDP session and the user data takes over. However, there are some parts of the network in which this rerouting mechanism does not protect against interruptions in service.

Pseudowire redundancy enables you to set up backup pseudowires. You can configure the network with redundant pseudowires and redundant network elements.

Prior to the failure of the primary pseudowire, the ability to switch traffic to the backup pseudowire is used to handle a planned pseudowire outage, such as router maintenance.

Configuration

This section describes the configuration for pseudowire redundancy.

```

/* Configure a cross-connect group with a static point-to-point
cross connect */
Router# configure
Router(config) # l2vpn
Router(config-l2vpn) # xconnect group A
Router(config-l2vpn-xc) # p2p xc1
Router(config-l2vpn-xc-p2p) # interface tengige 0/0/0/0.2
Router(config-l2vpn-xc-p2p) # neighbor 10.1.1.2 pw-id 2

```

```

/*Configure the pseudowire segment for the cross-connect group */
Router(config-l2vpn-xc-p2p-pw)#pw-class path1

/*Configure the backup pseudowire segment for the cross-connect group */
Router(config-l2vpn-xc-p2p-pw)# backup neighbor 10.2.2.2 pw-id 5
Router(config-l2vpn-xc-p2p-pw-backup)#end

/*Commit your configuration */
Router(config-l2vpn-xc-p2p-pw-backup)#commit
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]: yes

```

Running Configuration

```


Router# show-running configuration
...
l2vpn
encapsulation mpls
!
xconnect group A
p2p xc1
interface tengige 0/0/0/0.2
neighbor ipv4 10.1.1.2 pw-id 2
pw-class path1
backup neighbor 10.2.2.2 pw-id 5
!
!
...

```

Access Pseudowire Redundancy

The Access Pseudowire Redundancy feature allows you to configure a backup pseudowire under the bridge domain. When the primary pseudowire fails, the PE router switches to the backup pseudowire. The primary pseudowire resumes operation after it becomes functional. The primary pseudowire fails when the PE router fails or when there is a network outage.

Figure 11: Access Pseudowire Redundancy



Configure Access Pseudowire Redundancy

This section describes how you can configure access pseudowire redundancy.

Configuration Example

```
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# bridge group r1
Router(config-l2vpn-bg)# bridge-domain r1
Router(config-l2vpn-bg-bd)# interface TenGigE0/1/0/0.4
Router(config-l2vpn-bg-bd-ac)# exit
Router(config-l2vpn-bg-bd)# neighbor 10.0.0.1 pw-id 4
Router(config-l2vpn-bg-bd-pw)# backup neighbor 172.16.0.1 pw-id 4
Router(config-l2vpn-bg-bd-pw-backup)# commit
Router(config-l2vpn-bg-bd-pw-backup)# exit
```

```
Router# configure
Router(config)# interface TenGigE0/1/0/0.4 l2transport
Router(config-subif)# encapsulation dot1q 4
Router(config-subif)# rewrite ingress tag pop 1 symmetric
Router(config-subif)# commit
```

Running Configuration

This section shows access pseudowire redundancy running configuration.

```
configure
l2vpn
  bridge group r1
    bridge-domain r1
      interface TenGigE0/1/0/0.4
        !
        neighbor 10.0.0.1 pw-id 4
        backup neighbor 172.16.0.1 pw-id 4
        !
      !
    !
  !
interface TenGigE0/1/0/0.4 l2transport
  encapsulation dot1q 4
  rewrite ingress tag pop 1 symmetric
```

Verification

Verify the access pseudowire redundancy configuration.

```
Router# show l2vpn bridge-domain bd-name r1

Thu Apr 30 03:52:13.096 UTC
Legend: pp = Partially Programmed.
Bridge group: r1, bridge-domain: r1, id: 1, state: up, ShgId: 0, MSTi: 0
Aging: 300 s, MAC limit: 32000, Action: none, Notification: syslog
Filter MAC addresses: 0
ACs: 1 (1 up), VFIs: 0, PWs: 2 (1 up), PBBs: 0 (0 up), VNIs: 0 (0 up)
List of ACs:
  Te0/1/0/0.4, state: up, Static MAC addresses: 0
List of Access PWs:
  Neighbor 10.0.0.1 pw-id 4, state: up, Static MAC addresses: 0
  Neighbor 172.16.0.1 pw-id 4, state: standby, Static MAC addresses: 0, backup
List of VFIs:
List of Access VFIs:
```

Related Topics

- [Access Pseudowire Redundancy, on page 51](#)

Associated Commands

- show l2vpn bridge-domain

Virtual Circuit Connection Verification on L2VPN

Virtual Circuit Connection Verification (VCCV) is an L2VPN Operations, Administration, and Maintenance (OAM) feature that allows network operators to run IP-based provider edge-to-provider edge (PE-to-PE) keepalive protocol across a specified pseudowire to ensure that the pseudowire data path forwarding does not contain any faults. The disposition PE receives VCCV packets on a control channel, which is associated with the specified pseudowire. The control channel type and connectivity verification type, which are used for VCCV, are negotiated when the pseudowire is established between the PEs for each direction.

Two types of packets can arrive at the disposition egress:

- Type 1—Specifies normal Ethernet-over-MPLS (EoMPLS) data packets. This includes a) inband control word if negotiated during signalling and b) MPLS TTL expiry
- Type 2—Specifies a router alert label (label-0).

The router supports Label Switched Path (LSP) VCCV packets of Type 1. The VCCV echo reply is sent as an IPv4 packet, that is, the reply mode is IPv4.

The router does not support accounting of VCCV packets. .

GTP Load Balancing

The GPRS Tunneling Protocol (GTP) Load Balancing feature enables efficient distribution of traffic in mobile networks, and provides increased reliability and availability for the network.

GTP is a tunnel control and management protocol among General Packet Radio Service (GPRS) support nodes. Wireless networks use GTP tunnels to deliver mobile data. GTP includes GTP signaling (GTP-C) and data transfer (GTP-U) procedures. GTP-C specifies a tunnel control and management protocol, and is used to create, delete and modify tunnels. GTP-U uses a tunneling mechanism to provide a service for carrying user data packets over the network.

GTP load balancing is performed on IPv4 or IPv6 incoming packets with GTP payloads and on MPLS incoming labeled packets.


This feature supports GTP hashing only when the GTP UDP port is 2152.

The number of MPLS label stack in the transport layer is limited to three for GTP hashing. GTP hashing is not considered when the MPLS label stack exceeds three.

You need not reload the router after configuring or unconfiguring the **hw-module command** for GTP load balancing over MPLS to take effect.

The following figure shows an illustration of the mobile transport GTP-U load balancing.

Figure 12: Mobile Transport GTP-U Load-Balancing



The global L3 flow-based load balancing considers the following fields:

- source address
- destination address
- router ID
- source port
- destination port

For GTP traffic, however, the number of unique values for these fields is limited; this causes an uneven distribution of traffic. Sometimes, to facilitate redundancy and load balancing in a network, equal cost paths exist to different destinations. Load balancing does not occur in such scenarios as the source and destination IP addresses, as well as L4 ports, are the same. In order to achieve a greater distribution of traffic over equal cost links, load balancing (hashing) must occur on the GTP Tunnel Endpoint Identifier (TEID), which is unique for each traffic flow.

If the packet is UDP and the destination port is the GTP-U port (port number 2152), the GTP TEID is considered for load balancing. This provides GTP load balancing.

The TEID in the GTP header of a GTP packet identifies individual tunnel endpoints, thus achieving better mobile traffic load balancing within any given GRE tunnel. Additionally, this also helps in load balancing GTP traffic over Bundles at transit routers.

Load balancing based on tunnel endpoints is supported for Version 1 GTP packet and GTP version 2, if TEID is present. For GTP version 0, load balancing occurs in the same manner as before, as there is no TEID in version 0.



Note GTP load balancing is performed only for GTP-U (user data) packets. The GTP-C (control data) packets use a different destination port number of 2123 and hence, are subject to only the global L3 flow based load balancing.

By default, load balancing based on GTP-ID when GTP tunnel is over MPLS is disabled.

To enable GTP load balancing over MPLS, configure the **hw-module profile load-balance algorithm gtp-mpls** command.



CHAPTER 6

Configure Virtual LANs in Layer 2 VPNs

The Layer 2 Virtual Private Network (L2VPN) feature enables Service Providers (SPs) to provide L2 services to geographically disparate customer sites.

A virtual local area network (VLAN) is a group of devices on one or more LANs that are configured so that they can communicate as if they were attached to the same wire, when in fact they are located on a number of different LAN segments. The IEEE's 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames.

VLANs are very useful for user and host management, bandwidth allocation, and resource optimization. Using VLANs addresses the problem of breaking large networks into smaller parts so that broadcast and multicast traffic does not consume more bandwidth than necessary. VLANs also provide a higher level of security between segments of internal networks.

The 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames. Cisco IOS XR software supports VLAN sub-interface configuration on Gigabit Ethernet and 10-Gigabit Ethernet interfaces.

The configuration model for configuring VLAN Attachment Circuits (ACs) is similar to the model used for configuring basic VLANs, where the user first creates a VLAN sub-interface, and then configures that VLAN in sub-interface configuration mode. To create an Attachment Circuit, you need to include the **l2transport** keyword in the **interface** command string to specify that the interface is a L2 interface.

VLAN ACs support the following modes of L2VPN operation:

- **Basic Dot1Q Attachment Circuit**—The Attachment Circuit covers all frames that are received and sent with a specific VLAN tag.
- **QinQ Attachment Circuit**—The Attachment Circuit covers all frames received and sent with a specific outer VLAN tag and a specific inner VLAN tag. QinQ is an extension to Dot1Q that uses a stack of two tags.

Encapsulation

Encapsulation defines the matching criteria that maps a VLAN, a range of VLANs. Different types of encapsulations are default, dot1q, dot1ad. The following are the supported encapsulation types:

- **encapsulation default**: Configures the default service instance on a port.
- **encapsulation dot1q vlan extra-id** : Defines the matching criteria to map 802.1Q frames ingress on an interface to the appropriate service instance.

- **encapsulation dot1ad vlan-id** : Defines the matching criteria to map 802.1ad frames ingress on an interface to the appropriate service instance.
- **encapsulation dot1q second-dot1q**: Defines the matching criteria to map Q-in-Q ingress frames on an interface to the appropriate service instance.
- **encapsulation dot1ad dot1q**: Defines the matching criteria to be used in order to map single-tagged 802.1ad frames ingress on an interface to the appropriate service instance.

Restrictions and Limitations

To configure VLANs for Layer 2 VPNs, the following restrictions are applicable.

- In a point-to-point connection, the two Attachment Circuits do not have to be of the same type. For example, a port mode Ethernet Attachment Circuit can be connected to a Dot1Q Ethernet Attachment Circuit.
- Pseudowires can run in VLAN mode or in port mode. A pseudowire running in VLAN mode always carries Dot1Q or Dot1ad tag(s), while a pseudowire running in port mode may or may NOT carry tags. To connect these different types of circuits, popping, pushing, and rewriting tags is required.
- The Attachment Circuits on either side of an MPLS pseudowire can be of different types. In this case, the appropriate conversion is carried out at one or both ends of the Attachment Circuit to pseudowire connection.
- You can program a maximum number of 16 virtual MAC addresses on your router.
- [Configure VLAN Sub-Interfaces, on page 58](#)
- [Introduction to Ethernet Flow Point, on page 60](#)
- [Configure VLAN Header Rewrite, on page 65](#)

Configure VLAN Sub-Interfaces

Sub-interfaces are logical interfaces created on a hardware interface. These software-defined interfaces allow for segregation of traffic into separate logical channels on a single hardware interface as well as allowing for better utilization of the available bandwidth on the physical interface.

Sub-interfaces are distinguished from one another by adding an extension on the end of the interface name and designation. For instance, the Ethernet sub-interface 23 on the physical interface designated TenGigE 0/1/0/0 would be indicated by TenGigE 0/1/0/0.23.

Before a sub-interface is allowed to pass traffic, it must have a valid tagging protocol encapsulation and VLAN identifier assigned. All Ethernet sub-interfaces always default to the 802.1Q VLAN encapsulation. However, the VLAN identifier must be explicitly defined.

The sub-interface Maximum Transmission Unit (MTU) is inherited from the physical interface with 4 bytes allowed for the 802.1Q VLAN tag.

The following modes of VLAN sub-interface configuration are supported:

- Basic dot1q Attachment Circuit
- Q-in-Q Attachment Circuit

To configure a basic dot1q Attachment Circuit, use this encapsulation mode:

encapsulation dot1q *vlan extra-id*

To configure a basic dot1ad Attachment Circuit, use this encapsulation mode:

encapsulation dot1ad *vlan-id*

To configure a Q-in-Q Attachment Circuit, use the following encapsulation modes:

- **encapsulation dot1q** *vlan-id* **second-dot1q** *vlan-id*
- **encapsulation dot1ad** *vlan-id* **dot1q** *vlan-id*

Restrictions and Limitations

To configure VLAN sub-interface, the following restrictions are applicable.

- For double tagged packet, the VLAN range is supported only on the inner tag.
- VLANs separated by comma are called a VLAN lists. VLAN list are not supported on the router.
- If 0x9100/0x9200 is configured as tunneling ether-type, then dot1ad (0x88a8) encapsulation is not supported.
- If any sub-interface is already configured under a main interface, modifying the tunneling ether-type is not supported.
- You can program a maximum number of 16 virtual MAC addresses on your router.
- Following limitations are applicable to both outer and inner VLAN ranges:
 - 32 unique VLAN ranges are supported per system.
 - The overlap between outer VLAN ranges on sub-interfaces of the same Network Processor Unit (NPU) is not supported. A sub-interface with a single VLAN tag that falls into a range configured on another sub-interface of the same NPU is also considered an overlap.
 - The overlap between inner VLAN ranges on sub-interfaces of the same NPU is not supported.
 - Range 'any' does not result in explicit programming of a VLAN range in hardware and therefore does not count against the configured ranges.

Configuration Example

Configuring VLAN sub-interface involves:

- Creating a Ten Gigabit Ethernet sub-interface
- Enabling L2 transport mode on the interface
- Defining the matching criteria (encapsulation mode) to be used in order to map ingress frames on an interface to the appropriate service instance

Configuration of Basic dot1q Attachment Circuit**Running Configuration**

```
configure
```

```
interface TenGigE 0/0/0/10.1
  l2transport
  encapsulation dot1q 10 exact
!
```

Verification

Verify that the VLAN sub-interface is active:

```
router# show interfaces TenGigE 0/0/0/10.1

...
TenGigE0/0/0/10.1 is up, line protocol is up
  Interface state transitions: 1
  Hardware is VLAN sub-interface(s), address is 0011.1aac.a05a
  Layer 2 Transport Mode
  MTU 1518 bytes, BW 10000000 Kbit (Max: 10000000 Kbit)
    reliability Unknown, txload Unknown, rxload Unknown
  Encapsulation 802.1Q Virtual LAN,
    Outer Match: Dot1Q VLAN 10
    Ethertype Any, MAC Match src any, dest any
  loopback not set,
...

```

Associated Commands

- [encapsulation dot1ad dot1q](#)
- [encapsulation dot1q](#)
- [encapsulation dot1q second-dot1q](#)
- [l2transport \(Ethernet\)](#)
- [encapsulation dot1ad](#)

Introduction to Ethernet Flow Point

An Ethernet Flow Point (EFP) is a Layer 2 logical sub-interface used to classify traffic under a physical or a bundle interface. An EFP is defined by a set of filters (a set of entries) that are applied to all the ingress traffic to classify the frames that belong to a particular EFP. Each entry usually contains 0, 1 or 2 VLAN tags. You can specify a VLAN or QinQ tagging to match against on ingress. A packet that starts with the same tags as an entry in the filter is said to match the filter; if the start of the packet does not correspond to any entry in the filter, then the packet does not match the filter.

All traffic on ingress are processed by that EFP if a match occurs, and this can in turn change VLAN IDs, add or remove VLAN tags, and change ethertypes. After the frames are matched to a particular EFP, any appropriate feature (such as, any frame manipulations specified by the configuration as well as things such as QoS and ACLs) can be applied.

The benefits of EFP include:

- Identifying all frames that belong to a particular flow on a given interface

- Performing VLAN header rewrites
(See, [Configure VLAN Header Rewrite, on page 65](#))
- Adding features to the identified frames
- Optionally defining how to forward the identified frames in the data path

Limitations of EFP

Egress EFP filtering is not supported on Cisco IOS XR.

Identify Frames of an EFP

The EFP identifies frames belonging to a particular flow on a given port, independent of their Ethernet encapsulation. An EFP can flexibly map frames into a flow or EFP based on the fields in the frame header. The frames can be matched to an EFP using VLAN tag(s).

The frames cannot be matched to an EFP through this:

- Any information outside the outermost Ethernet frame header and its associated tags such as
 - IPv4, IPv6, or MPLS tag header data
 - C-DMAC, C-SMAC, or C-VLAN

VLAN Tag Identification

Below table describes the different encapsulation types and the EFP identifier corresponding to each.

Encapsulation Type	EFP Identifier
Single tagged frames	802.1Q customer-tagged Ethernet frames
Double tagged frames	802.1Q (ethertype 0x9100) double tagged frames 802.1ad (ethertype 0x9200) double tagged frames

You can use wildcards while defining frames that map to a given EFP. EFPs can distinguish flows based on a single VLAN tag, a stack of VLAN tags or a combination of both (VLAN stack with wildcards). It provides the EFP model, a flexibility of being encapsulation agnostic, and allows it to be extensible as new tagging or tunneling schemes are added.

Apply Features

After the frames are matched to a particular EFP, any appropriate features can be applied. In this context, “features” means any frame manipulations specified by the configuration as well as things such as QoS and ACLs. The Ethernet infrastructure provides an appropriate interface to allow the feature owners to apply their features to an EFP. Hence, IM interface handles are used to represent EFPs, allowing feature owners to manage their features on EFPs in the same way the features are managed on regular interfaces or sub-interfaces.

The only L2 features that can be applied on an EFP that is part of the Ethernet infrastructure are the L2 header encapsulation modifications. The L2 features are described in this section.

Encapsulation Modifications

EFP supports these L2 header encapsulation modifications on both ingress and egress:

- Push 1 or 2 VLAN tags
- Pop 1 or 2 VLAN tags



Note This modification can only pop tags that are matched as part of the EFP.

- Rewrite 1 or 2 VLAN tags:
 - Rewrite outer tag
 - Rewrite outer 2 tags
 - Rewrite outer tag and push an additional tag

For each of the VLAN ID manipulations, these can be specified:

- The VLAN tag type, that is, C-VLAN, S-VLAN, or I-TAG. The ethertype of the 802.1Q C-VLAN tag is defined by the dot1q tunneling type command.
- The VLAN ID. 0 can be specified for an outer VLAN tag to generate a priority-tagged frame.



Note For tag rewrites, the CoS bits from the previous tag should be preserved in the same way as the DEI bit for 802.1ad encapsulated frames.

Define Data-Forwarding Behavior

The EFP can be used to designate the frames belonging to a particular Ethernet flow forwarded in the data path. These forwarding cases are supported for EFPs in Cisco IOS XR software:

- L2 Switched Service (Bridging)—The EFP is mapped to a bridge domain, where frames are switched based on their destination MAC address. This includes multipoint services:
 - Ethernet to Ethernet Bridging
 - Multipoint Layer 2 Services
- L2 Stitched Service (AC to AC xconnect)—This covers point-to-point L2 associations that are statically established and do not require a MAC address lookup.
 - Ethernet to Ethernet Local Switching—The EFP is mapped to an S-VLAN either on the same port or on another port. The S-VLANs can be identical or different.
- Tunneled Service (xconnect)—The EFP is mapped to a Layer 3 tunnel. This covers point-to-point services, such as EoMPLS.

Ethernet Flow Points Visibility

EFP Visibility feature enables you to configure multiple VLANs in the same bridge-domain.

An Ethernet flow point (EFP) service instance is a logical interface that connects a bridge domain to a physical port or to an EtherChannel group. A VLAN tag is used to identify the EFP.

Earlier only one EFP was allowed per bridge-domain. With EFP visibility feature, you can configure a maximum of:

- 600 EFPs per bridge-domain.
- 100 EFPs per port.

Irrespective of number of ports available, you have flexibility to add more EFPs in one bridge group.

Configuring EFP Visibility

This example shows how to configure IGMP snooping on VLAN interfaces under a bridge domain with multiple EFPs.

```

/* Configure two IGMP Snooping profiles */
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# igmp snooping profile 1
RP/0/RP0/CPU0:router(config-igmp-snooping-profile)# exit
RP/0/RP0/CPU0:router(config)# igmp snooping profile 2
RP/0/RP0/CPU0:router(config-igmp-snooping-profile)#commit

!

/* Configure VLAN interfaces for L2 transport */
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface gigabitEthernet 0/8/0/8
RP/0/RP0/CPU0:router(config-if)# bundle id 2 mode on
RP/0/RP0/CPU0:router(config-if)# no shut
RP/0/RP0/CPU0:router(config-if)# exit
RP/0/RP0/CPU0:router(config)# interface gigabitEthernet 0/8/0/9
RP/0/RP0/CPU0:router(config-if)# bundle id 3 mode on
RP/0/RP0/CPU0:router(config-if)# no shut
RP/0/RP0/CPU0:router(config-if)# exit

RP/0/RP0/CPU0:router(config)# interface Bundle-Ether2
RP/0/RP0/CPU0:router(config-if)# exit
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether3
RP/0/RP0/CPU0:router(config-if)# exit

RP/0/RP0/CPU0:router(config)# interface Bundle-Ether2.2 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 2
RP/0/RP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether2.3 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 3
RP/0/RP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether2.4 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 4
RP/0/RP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether2.5 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 5
RP/0/RP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
RP/0/RP0/CPU0:router(config-subif)# exit

RP/0/RP0/CPU0:router(config)# interface Bundle-Ether3.2 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 2
RP/0/RP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
RP/0/RP0/CPU0:router(config-subif)# exit

```

```

RP/0/RP0/CPU0:router(config)# interface Bundle-Ether3.3 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 3
RP/0/RP0/CPU0:router(config-subif)# rewrite ingress tag pop 1 symmetric
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# commit

/* Attach a profile and add interfaces to the bridge domain.
Attach a profile to one of the interfaces. The other interface
inherits IGMP snooping configuration attributes from the bridge domain profile */

RP/0/RP0/CPU0:router(config)#l2vpn
RP/0/RP0/CPU0:router(config-l2vpn)#bridge group VLAN2
RP/0/RP0/CPU0:router(config-l2vpn-bg)#bridge-domain VLAN2
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#efp-visibility
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#igmp snooping profile 1
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface bundle-Ether2.2
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface bundle-Ether 2.3
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface bundle-Ether 2.4
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface bundle-Ether 2.5
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg)#bridge-domain vlan3
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#efp-visibility
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#igmp snooping profile 2
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface bundle-Ether3.2
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#interface bundle-Ether 3.3
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#routed interface bvi2
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-bvi)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#evi 2
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd-evi)#exit
RP/0/RP0/CPU0:router(config-l2vpn-bg-bd)#commit

```

Verification

Verify the configured bridge ports:

```

RP/0/RP0/CPU0:router# show igmp snooping port
Bridge Domain VLAN2:VLAN2

```

Port	State				
	Oper	STP	Red	#Grps	#SGs
----	----	---	---	-----	-----
BVI2	Up	-	-	0	0
Bundle-Ether2.2	Up	-	-	100	0
Bundle-Ether2.3	Up	-	-	100	0
Bundle-Ether2.4	Up	-	-	100	0
Bundle-Ether2.5	Up	-	-	100	0

```

Bridge Domain VLAN3:VLAN3

```

Port	State				
	Oper	STP	Red	#Grps	#SGs
----	----	---	---	-----	-----
BVI3	Up	-	-	0	0
Bundle-Ether3.2	Up	-	-	100	0
Bundle-Ether3.3	Up	-	-	100	0

In the above output verify the status of BVI and EFPs are **Up**, and the **#Grps** and **#SG** show the correct number of IGMP join received.

Configure VLAN Header Rewrite

EFP supports the following VLAN header rewrites on both ingress and egress ports:

- Push 1 VLAN tag
- Pop 1 VLAN tag



Note This rewrite can only pop tags that are matched as part of the EFP.

- Translate 1 or 2 VLAN tags:
 - Translate 1-to-1 tag: Translates the outermost tag to another tag
 - Translate 1-to-2 tags: Translates the outermost tag to two tags
 - Translate 2-to-2 tags: Translates the outermost two tags to two other tags

Various combinations of ingress, egress VLAN rewrites with corresponding tag actions during ingress and egress VLAN translation, are listed in the following sections:

Limitations

The limitations for VLAN header rewrites are as follows:

- Push 1 is not supported for dot1ad configuration.
- Push 2 is supported only on:
 - Untagged EFP
 - Dot1q EFP with **exact** configuration statement
- Translate 1 to 1 is not supported for dot1ad configuration.
- Translate 1 to 2 is not supported with **dot1q tunneling ethertype** configuration statement.
- Pop 2 is not supported.
- Translate 2 to 1 is not supported.
- When a single-tag range is used, double tagged traffic does not match.

For example, in the following configuration, dot1q 2-6 is the outer tag.

```
Router#configure
Router(config)# interface GigabitEthernet0/0/0/0.0 12transport
Router(config-if)# encapsulation dot1q 2-6
```

- An incoming packet with an outer tag of 2 and **ANY** inner tag does not match. For example, the double tag packet of outer tag 2 and inner tag 1 is not be accepted on the interface 0/0/0/0.0.

- But, an incoming packet with a single tag of 2 is accepted. For example, the single tag packet of outer tag between 2 to 6 is accepted on the interface 0/0/0/0.0.

Configuration Example

This topic covers VLAN header rewrites on various attachment circuits, such as:

- L2 single-tagged sub-interface
- L2 double-tagged sub-interface

Configuring VLAN header rewrite involves:

- Creating a TenGigabit Ethernet sub-interface
- Enabling L2 transport mode on the interface
- Defining the matching criteria (encapsulation mode) to be used in order to map single-tagged frames ingress on an interface to the appropriate service instance
- Specifying the encapsulation adjustment that is to be performed on the ingress frame

Configuration of VLAN Header Rewrite (single-tagged sub-interface)

```
Router# configure
Router(config)# interface TenGigE 0/0/0/10.1 l2transport
Router(config-if)# encapsulation dot1q 10 exact
Router(config-if)# rewrite ingress tag push dot1q 20 symmetric
```

Running Configuration

```
/* Configuration without rewrite */

configure
interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10 exact
!
!

/* Configuration with rewrite */

/* PUSH 1 */
interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10
  rewrite ingress tag push dot1q 20 symmetric
!
!

/* POP 1 */
interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10
  rewrite ingress tag pop 1
!
!

/* TRANSLATE 1-1 */
```

```

interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10
  rewrite ingress tag translate 1-to-1 dot1q 20
  !
!

/* TRANSLATE 1-2 */

interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10
  rewrite ingress tag translate 1-to-2 dot1q 20 second-dot1q 30
  !
!

```

Running Configuration (VLAN header rewrite on double-tagged sub-interface)

```

/* Configuration without rewrite */

interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10 second-dot1q 11
  !
!

/* Configuration with rewrite */

/* PUSH 1 */
interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10 second-dot1q 11
  rewrite ingress tag push dot1q 20 symmetric
  !
!

/* TRANSLATE 1-1 */

interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10 second-dot1q 11
  rewrite ingress tag translate 1-to-1 dot1q 20
  !
!

/* TRANSLATE 1-2 */

interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10 second-dot1q 11
  rewrite ingress tag translate 1-to-2 dot1q 20 second-dot1q 30
  !
!

/* TRANSLATE 2-2 */

interface TenGigE0/0/0/0.1 l2transport
  encapsulation dot1q 10 second-dot1q 11
  rewrite ingress tag translate 2-to-2 dot1q 20 second-dot1q 30
  !
!

```

Associated Commands

- [encapsulation dot1ad dot1q](#)

- [encapsulation dot1q](#)
- [encapsulation dot1q second-dot1q](#)
- [l2transport \(Ethernet\)](#)
- [rewrite ingress tag](#)

Valid Ingress Rewrite Actions

Table 1: Valid Ingress Rewrite Actions

Interface Configuration	Ingress Rewrite Action
dot1q	No rewrite
dot1q	Pop 1
dot1q	Push 1
dot1q	Push 2
dot1q	Translate 1 to 1
dot1q	Translate 1 to 2
QinQ	No rewrite
QinQ	Pop 1
QinQ	Push 1
QinQ	Translate 1 to 1
QinQ	Translate 1 to 2
QinQ	Translate 2 to 2
Untagged	No rewrite
Untagged	Push 1
Untagged	Push 2

The following notations are used for the rewrite actions mentioned in the table:

- Translate 1-to-1 tag: Translates the outermost tag to another tag.
- Translate 1-to-2 tags: Translates the outermost tag to two tags.
- Translate 2-to-2 tags: Translates the outermost two tags to two other tags.

Valid Ingress-Egress Rewrite Combinations

Table 2: Valid Ingress-Egress Rewrite Combinations

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
dot1q	No rewrite	dot1q	No rewrite
dot1q	No rewrite	dot1q	Pop 1
dot1q	No rewrite	dot1q	Push 1
dot1q	No rewrite	dot1q	Translate 1-to-1
dot1q	Pop 1	dot1q	No rewrite
dot1q	Pop 1	dot1q	Pop 1
dot1q	Push 1	dot1q	No rewrite
dot1q	Push 1	dot1q	Push 1
dot1q	Push 1	dot1q	Push 2
dot1q	Push 1	dot1q	Translate 1-to-1
dot1q	Push 1	dot1q	Translate 1-to-2
dot1q	Push 2 / Translate 1-to-2	dot1q	Push 1
dot1q	Push 2 / Translate 1-to-2	dot1q	Push 2
dot1q	Push 2 / Translate 1-to-2	dot1q	Translate 1-to-2
dot1q	Translate 1-to-1	dot1q	No rewrite
dot1q	Translate 1-to-1	dot1q	Push 1
dot1q	Translate 1-to-1	dot1q	Translate 1-to-1
dot1q	No rewrite	dot1q range	No rewrite
dot1q	No rewrite	dot1q range	Push 1
dot1q	Pop 1	dot1q range	No rewrite
dot1q	Push 1	dot1q range	No rewrite
dot1q	Push 1	dot1q range	Push 1
dot1q	Push 1	dot1q range	Push 2
dot1q	Translate 1-to-1	dot1q range	No rewrite

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
dot1q	Translate 1-to-1	dot1q range	Push 1
dot1q	Translate 1-to-2	dot1q range	Push 1
dot1q	Translate 1-to-2	dot1q range	Push 2
dot1q	No rewrite / Translate 1-to-1	QinQ	No rewrite
dot1q	No rewrite / Translate 1-to-1	QinQ	Pop 1
dot1q	No rewrite / Translate 1-to-1	QinQ	Push 1
dot1q	No rewrite / Translate 1-to-1	QinQ	Translate 1-to-1
dot1q	Pop 1	QinQ	No rewrite
dot1q	Pop 1	QinQ	Pop 1
dot1q	Push 1	QinQ	No rewrite
dot1q	Push 1	QinQ	Pop 1
dot1q	Push 1	QinQ	Push 1
dot1q	Push 1	QinQ	Translate 1-to-1
dot1q	Push 1	QinQ	Translate 1-to-2
dot1q	Push 1	QinQ	Translate 2-to-2
dot1q	Push 2 / Translate 1-to-2	QinQ	No rewrite
dot1q	Push 2 / Translate 1-to-2	QinQ	Push 1
dot1q	Push 2 / Translate 1-to-2	QinQ	Translate 1-to-1
dot1q	Push 2 / Translate 1-to-2	QinQ	Translate 1-to-2
dot1q	Push 2 / Translate 1-to-2	QinQ	Translate 2-to-2
dot1q	No rewrite / Translate 1-to-1	QinQ range	No rewrite
dot1q	No rewrite / Translate 1-to-1	QinQ range	Pop 1
dot1q	No rewrite / Translate 1-to-1	QinQ range	Push 1
dot1q	No rewrite / Translate 1-to-1	QinQ range	Translate 1-to-1
dot1q	Pop 1	QinQ range	No rewrite
dot1q	Pop 1	QinQ range	Pop 1
dot1q	Push 1	QinQ range	No rewrite

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
dot1q	Push 1	QinQ range	Pop 1
dot1q	Push 1	QinQ range	Push 1
dot1q	Push 1	QinQ range	Translate 1-to-1
dot1q	Push 1	QinQ range	Translate 1-to-2
dot1q	Push 2 / Translate 1-to-2	QinQ range	No rewrite
dot1q	Push 2 / Translate 1-to-2	QinQ range	Push 1
dot1q	Push 2 / Translate 1-to-2	QinQ range	Translate 1-to-1
dot1q	Push 2 / Translate 1-to-2	QinQ range	Translate 1-to-2
dot1q	No rewrite	QinQ range	No rewrite
dot1q	No rewrite	Untagged	Push 1
dot1q	Pop 1	Untagged	No rewrite
dot1q	Push 1	Untagged	Push 1
dot1q	Push 1	Untagged	Push 2
dot1q	Push 2	Untagged	Push 2
dot1q	Translate 1-to-1	Untagged	Push 1
dot1q	Translate 1-to-2	Untagged	Push 2
dot1q range	No rewrite	dot1q range	No rewrite
dot1q range	No rewrite	dot1q range	Push 1
dot1q range	Push 1	dot1q range	No rewrite
dot1q range	Push 1	dot1q range	Push 1
dot1q range	Push 1	dot1q range	Push 2
dot1q range	Push 2	dot1q range	Push 1
dot1q range	Push 2	dot1q range	Push 2
dot1q range	No rewrite	QinQ	No rewrite
dot1q range	No rewrite	QinQ	Pop 1
dot1q range	No rewrite	QinQ	Push 1
dot1q range	No rewrite	QinQ	Translate 1-to-1

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
dot1q range	Push 1	QinQ	No rewrite
dot1q range	Push 1	QinQ	Pop 1
dot1q range	Push 1	QinQ	Push 1
dot1q range	Push 1	QinQ	Translate 1-to-1
dot1q range	Push 1	QinQ	Translate 1-to-2
dot1q range	Push 1	QinQ	Translate 2-to-2
dot1q range	Push 2	QinQ	No rewrite
dot1q range	Push 2	QinQ	Push 1
dot1q range	Push 2	QinQ	Translate 1-to-1
dot1q range	Push 2	QinQ	Translate 1-to-2
dot1q range	Push 2	QinQ	Translate 2-to-2
dot1q range	No rewrite	QinQ range / QinAny	No rewrite
dot1q range	No rewrite	QinQ range	Pop 1
dot1q range	No rewrite	QinQ range / QinAny	Push 1
dot1q range	No rewrite	QinQ range / QinAny	Translate 1-to-1
dot1q range	Push 1	QinQ range / QinAny	No rewrite
dot1q range	Push 1	QinQ range / QinAny	Pop 1
dot1q range	Push 1	QinQ range	Push 1
dot1q range	Push 1	QinQ range / QinAny	Translate 1-to-1
dot1q range	Push 1	QinQ range / QinAny	Translate 1-to-2

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
dot1q range	Push 2	QinQ range / QinAny	No rewrite
dot1q range	Push 2	QinQ range / QinAny	Push 1
dot1q range	Push 2	QinQ range / QinAny	Translate 1-to-1
dot1q range	Push 2	QinQ range / QinAny	Translate 1-to-2
dot1q range	No rewrite	Untagged	No rewrite
dot1q range	No rewrite	Untagged	Push 1
dot1q range	Push 1	Untagged	Push 1
dot1q range	Push 1	Untagged	Push 2
dot1q range	Push 2	Untagged	Push 2
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ	No rewrite
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ	Pop 1
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ	Push 1
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ	Translate 1-to-1
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ	Translate 1-to-2
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ	Translate 2-to-2
QinQ	Pop 1	QinQ	No rewrite
QinQ	Pop 1	QinQ	Pop 1
QinQ	Pop 1	QinQ	Push 1
QinQ	Pop 1	QinQ	Translate 1-to-1
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ	No rewrite

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ	Push 1
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ	Translate 1-to-1
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ	Translate 1-to-2
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ	Translate 2-to-2
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	No rewrite
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Pop 1
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Push 1
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Translate 1-to-1
QinQ	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Translate 1-to-2
QinQ	Pop 1	QinQ range / QinAny	No rewrite
QinQ	Pop 1	QinQ range / QinAny	Pop 1
QinQ	Pop 1	QinQ range / QinAny	Push 1
QinQ	Pop 1	QinQ range / QinAny	Translate 1-to-1
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ range / QinAny	No rewrite
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ range / QinAny	Push 1
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ range / QinAny	Translate 1-to-1

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
QinQ	Translate 1-to-2 / Translate 2-to-2	QinQ range / QinAny	Translate 1-to-2
QinQ	No rewrite	Untagged	No rewrite
QinQ	No rewrite	Untagged	Push 1
QinQ	No rewrite	Untagged	Push 2
QinQ	Pop 1	Untagged	No rewrite
QinQ	Pop 1	Untagged	Push 1
QinQ	Push 1 / Translate 1-to-1	Untagged	Push 1
QinQ	Push 1 / Translate 1-to-1	Untagged	Push 2
QinQ	Translate 1-to-2 / Translate 2-to-2	Untagged	Push 2
QinQ range / QinAny	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	No rewrite
QinQ range / QinAny	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Pop 1
QinQ range / QinAny	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Push 1
QinQ range / QinAny	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Translate 1-to-1
QinQ range / QinAny	No rewrite / push 1 / Translate 1-to-1	QinQ range / QinAny	Translate 1-to-2
QinQ range / QinAny	Pop 1	QinQ range / QinAny	No rewrite
QinQ range / QinAny	Pop 1	QinQ range / QinAny	Pop 1
QinQ range / QinAny	Pop 1	QinQ range / QinAny	Push 1
QinQ range	Pop 1	QinQ range	Translate 1-to-1

Ingress Interface Configuration	Ingress Interface Rewrite Action	Egress Interface Configuration	Egress Interface Rewrite Action
QinQ range / QinAny	Translate 1-to-2	QinQ range / QinAny	No rewrite
QinQ range / QinAny	Translate 1-to-2	QinQ range / QinAny	Push 1
QinQ range / QinAny	Translate 1-to-2	QinQ range / QinAny	Translate 1-to-1
QinQ range / QinAny	Translate 1-to-2	QinQ range / QinAny	Translate 1-to-2
QinQ range / QinAny	No rewrite	Untagged	No rewrite
QinQ range / QinAny	No rewrite	Untagged	Push 1
QinQ range / QinAny	No rewrite	Untagged	Push 2
QinQ range / QinAny	Pop 1	Untagged	No rewrite
QinQ range / QinAny	Pop 1	Untagged	Push 1
QinQ range / QinAny	Push 1 / Translate 1-to-1	Untagged	Push 1
QinQ range / QinAny	Push 1 / Translate 1-to-1	Untagged	Push 2
QinQ range / QinAny	Translate 1-to-2	Untagged	Push 2
Untagged	No rewrite	Untagged	No rewrite
Untagged	Push 1	Untagged	Push 1
Untagged	Push 2	Untagged	Push 2

The following notations are used for the rewrite actions mentioned in the table:

- Translate 1-to-1 tag: Translates the outermost tag to another tag

- Translate 1-to-2 tags: Translates the outermost tag to two tags
- Translate 2-to-2 tags: Translates the outermost two tags to two other tags



CHAPTER 7

EVPN Features

This chapter describes how to configure Layer 2 Ethernet VPN (EVPN) features on the router.

- [EVPN Overview, on page 79](#)
- [EVPN Concepts, on page 80](#)
- [EVPN Operation, on page 81](#)
- [EVPN Route Types, on page 82](#)
- [Configure EVPN L2 Bridging Service, on page 83](#)
- [EVPN Software MAC Learning , on page 84](#)
- [EVPN Out of Service, on page 92](#)
- [EVPN Multiple Services per Ethernet Segment, on page 95](#)
- [Network Convergence using Core Isolation Protection, on page 101](#)
- [Conditional Advertisement of Default-Originate, on page 107](#)
- [EVPN Routing Policy, on page 111](#)
- [Support for DHCPv4 and DHCPv6 Client over BVI, on page 126](#)

EVPN Overview

Ethernet VPN (EVPN) is a next generation solution that provides Ethernet multipoint services over MPLS networks. EVPN operates in contrast to the existing Virtual Private LAN Service (VPLS) by enabling control-plane based MAC learning in the core. In EVPN, PEs participating in the EVPN instances learn customer MAC routes in control-plane using MP-BGP protocol. Control-plane MAC learning brings a number of benefits that allow EVPN to address the VPLS shortcomings, including support for multi-homing with per-flow load balancing.

EVPN provides the solution for network operators for the following emerging needs in their network:

- Data center interconnect operation (DCI)
- Cloud and services virtualization
- Remove protocols and network simplification
- Integration of L2 and L3 services over the same VPN
- Flexible service and workload placement
- Multi-tenancy with L2 and L3 VPN
- Optimal forwarding and workload mobility

- Fast convergence
- Efficient bandwidth utilization

EVPN Benefits

The EVPN provides the following benefits:

- **Integrated Services:** Integrated L2 and L3 VPN services, L3VPN-like principles and operational experience for scalability and control, all-active multi-homing and PE load-balancing using ECMP, and enables load balancing of traffic to and from CEs that are multihomed to multiple PEs.
- **Network Efficiency:** Eliminates flood and learn mechanism, fast-reroute, resiliency, and faster reconvergence when the link to dual-homed server fails, optimized Broadcast, Unknown-unicast, Multicast (BUM) traffic delivery.
- **Service Flexibility:** MPLS data plane encapsulation, support existing and new services types (E-LAN, E-Line), peer PE auto-discovery, and redundancy group auto-sensing.

EVPN Modes

The following EVPN modes are supported:

- **Single-homing** - This enables you to connect a customer edge (CE) device to one provider edge (PE) device.
- **Multihoming** - This enables you to connect a customer edge (CE) device to more than one provider edge (PE) device. Multihoming ensures redundant connectivity. The redundant PE device ensures that there is no traffic disruption when there is a network failure. Following are the types of multihoming:
 - **All-Active** - In all-active mode all the PEs attached to the particular Ethernet-Segment is allowed to forward traffic to and from that Ethernet Segment.

EVPN Concepts

To implement EVPN features, you need to understand the following concepts:

- **Ethernet Segment (ES):** An Ethernet segment is a set of Ethernet links that connects a multihomed device. If a multi-homed device or network is connected to two or more PEs through a set of Ethernet links, then that set of links is referred to as an Ethernet segment. The Ethernet segment route is also referred to as Route Type 4. This route is used for designated forwarder (DF) election for BUM traffic.
- **Ethernet Segment Identifier (ESI):** Ethernet segments are assigned a unique non-zero identifier, which is called an Ethernet Segment Identifier (ESI). ESI represents each Ethernet segment uniquely across the network.
- **EVI:** The EVPN instance (EVI) is represented by the virtual network identifier (VNI). An EVI represents a VPN on a PE router. It serves the same role of an IP VPN Routing and Forwarding (VRF), and EVIs are assigned import/export Route Targets (RTs). Depending on the service multiplexing behaviors at the User to Network Interface (UNI), all traffic on a port (all-to-one bundling), or traffic on a VLAN (one-to-one mapping), or traffic on a list/range of VLANs (selective bundling) can be mapped to a Bridge Domain (BD). This BD is then associated to an EVI for forwarding towards the MPLS core.


- **EAD/ES:** Ethernet Auto Discovery Route per ES is also referred to as Route Type 1. This route is used to converge the traffic faster during access failure scenarios. This route has Ethernet Tag of 0xFFFFFFFF.
- **EAD/EVI:** Ethernet Auto Discovery Route per EVI is also referred to as Route Type 1. This route is used for aliasing and load balancing when the traffic only hashes to one of the switches. This route cannot have Ethernet tag value of 0xFFFFFFFF to differentiate it from the EAD/ES route.
- **Aliasing:** It is used for load balancing the traffic to all the connected switches for a given Ethernet segment using the Route Type 1 EAD/EVI route. This is done irrespective of the switch where the hosts are actually learned.
- **Mass Withdrawal:** It is used for fast convergence during the access failure scenarios using the Route Type 1 EAD/ES route.
- **DF Election:** It is used to prevent forwarding of the loops. Only a single router is allowed to decapsulate and forward the traffic for a given Ethernet Segment.

EVPN Operation

At startup, PEs exchange EVPN routes in order to advertise the following:

- **VPN membership:** The PE discovers all remote PE members of a given EVI. In the case of a multicast ingress replication model, this information is used to build the PEs flood list associated with an EVI. BUM labels and unicast labels are exchanged when MAC addresses are learned.
- **Ethernet segment reachability:** In multihoming scenarios, the PE auto-discovers remote PE and their corresponding redundancy mode (all-active or single-active). In case of segment failures, PEs withdraw the routes used at this stage in order to trigger fast convergence by signaling a MAC mass withdrawal on remote PEs.
- **Redundancy Group membership:** PEs connected to the same Ethernet segment (multihoming) automatically discover each other and elect a Designated Forwarder (DF) that is responsible for forwarding Broadcast, Unknown unicast and Multicast (BUM) traffic for a given EVI.

Figure 13: EVPN Operation



EVPN can operate in single-homing or dual-homing mode. Consider single-homing scenario, when EVPN is enabled on PE, Route Type 3 is advertised where each PE discovers all other member PEs for a given EVPN

instance. When an unknown unicast (or BUM) MAC is received on the PE, it is advertised as EVPN Route Type 2 to other PEs. MAC routes are advertised to the other PEs using EVPN Route Type 2. In multihoming scenarios, Route Types 1, 3, and 4 are advertised to discover other PEs and their redundancy modes (single-active or all-active). Use of Route Type 1 is to auto-discover other PE which hosts the same CE. The other use of this route type is to fast route unicast traffic away from a broken link between CE and PE. Route Type 4 is used for electing designated forwarder. For instance, consider the topology when customer traffic arrives at the PE, EVPN MAC advertisement routes distribute reachability information over the core for each customer MAC address learned on local Ethernet segments. Each EVPN MAC route announces the customer MAC address and the Ethernet segment associated with the port where the MAC was learned from and its associated MPLS label. This EVPN MPLS label is used later by remote PEs when sending traffic destined to the advertised MAC address.

EVPN Route Types

The EVPN network layer reachability information (NLRI) provides different route types.

Table 3: EVPN Route Types

Route Type	Name	Usage
1	Ethernet Auto-Discovery (AD) Route	Few routes are sent per ES, carries the list of EVIs that belong to ES
2	MAC/IP Advertisement Route	Advertise MAC, address reachability, advertise IP/MAC binding
3	Inclusive Multicast Ethernet Tag Route	Multicast Tunnel End point discovery
4	Ethernet Segment Route	Redundancy group discovery, DF election
5	IP Prefix Route	Advertise IP prefixes.

Route Type 1: Ethernet Auto-Discovery (AD) Route

The Ethernet Auto-Discovery (AD) routes are advertised on per EVI and per ESI basis. These routes are sent per ES. They carry the list of EVIs that belong to the ES. The ESI field is set to zero when a CE is single-homed. This route type is used for mass withdrawal of MAC addresses and aliasing for load balancing.

Route Type 2: MAC/IP Advertisement Route

These routes are per-VLAN routes, so only PEs that are part of a VNI require these routes. The host's IP and MAC addresses are advertised to the peers within NRLI. The control plane learning of MAC addresses reduces unknown unicast flooding.

Route Type 3: Inclusive Multicast Ethernet Tag Route

This route establishes the connection for broadcast, unknown unicast, and multicast (BUM) traffic from a source PE to a remote PE. This route is advertised on per VLAN and per ESI basis.

Route Type 4: Ethernet Segment Route

Ethernet segment routes enable to connect a CE device to two or PE devices. ES route enables the discovery of connected PE devices that are connected to the same Ethernet segment.

Route Type 5: IP Prefix Route

The IP prefixes are advertised independently of the MAC-advertised routes. With EVPN IRB, host route /32 is advertised using RT-2 and subnet /24 is advertised using RT-5.



Note With EVPN IRB, host route /32 are advertised using RT-2 and subnet /24 are advertised using RT-5.

Configure EVPN L2 Bridging Service

Perform the following steps to configure EVPN L2 bridging service.



Note Always ensure to change the label mode from per-prefix to per-VRF label mode. Since L2FIB and VPNv4 route (labels) shares the same resource, BVI ping fails when you exhaust the resources.



Note Flooding disable isn't supported on EVPN bridge domains.

```

/* Configure address family session in BGP */
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router# (config)# router bgp 200
RP/0/RSP0/CPU0:router# (config-bgp)# bgp router-id 209.165.200.227
RP/0/RSP0/CPU0:router# (config-bgp)# address-family l2vpn evpn
RP/0/RSP0/CPU0:router# (config-bgp)# neighbor 10.10.10.10
RP/0/RSP0/CPU0:router# (config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router# (config-bgp-nbr)# description MPLSFACING-PEER
RP/0/RSP0/CPU0:router# (config-bgp-nbr)# update-source Loopback 0
RP/0/RSP0/CPU0:router# (config-bgp-nbr)# address-family l2vpn evpn

/* Configure EVI and define the corresponding BGP route targets */

Router# configure
Router(config)# evpn
Router(config-evpn)# evi 6005
Router(config-evpn-evi)# bgp
Router(config-evpn-evi-bgp)# rd 200:50
Router(config-evpn-evi-bgp)# route-target import 100:6005
Router(config-evpn-evi-bgp)# route-target export 100:6005
Router(config-evpn-evi-bgp)# exit
Router(config-evpn-evi)# advertise-mac

/* Configure a bridge domain */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# bridge group 1

```

```

Router(config-l2vpn-bg)# bridge-domain 1-1
Router(config-l2vpn-bg-bd)# interface GigabitEthernet 0/0/0/1
Router(config-l2vpn-bg-bd-ac)# evi 6005
Router(config-l2vpn-bg-bd-ac-evi)# commit
Router(config-l2vpnbg-bd-ac-evi)# exit

```

Running Configuration

```

router bgp 200 bgp
  router-id 209.165.200.227
  address-family l2vpn evpn
  neighbor 10.10.10.10
    remote-as 200 description MPLS-FACING-PEER
    updatesource Loopback0
    addressfamily l2vpn evpn
  !

configure
  evpn
  evi 6005
  bgp
    rd 200:50
    route-target import 100:6005
    route-target export 100:6005
  !
  advertise-mac

configure
  l2vpn
  bridge group 1
  bridge-domain 1-1
    interface GigabitEthernet 0/0/0/1


    evi 6005
  !

```

EVPN Software MAC Learning

The MAC addresses learned on one device needs to be learned or distributed on the other devices in a VLAN. EVPN Software MAC Learning feature enables the distribution of the MAC addresses learned on one device to the other devices connected to a network. The MAC addresses are learnt from the remote devices using BGP.

Figure 14: EVPN Software MAC Learning



The above figure illustrates the process of software MAC learning. The following are the steps involved in the process:

1. Traffic comes in on one port in the bridge domain.
2. The source MAC address (AA) is learnt on the PE and is stored as a dynamic MAC entry.
3. The MAC address (AA) is converted into a type-2 BGP route and is sent over BGP to all the remote PEs in the same EVI.
4. The MAC address (AA) is updated on the PE as a remote MAC address.

Configure EVPN Software MAC Learning

The following section describes how you can configure EVPN Software MAC Learning:



Note The router does not support flow-aware transport (FAT) pseudowire.

```

/* Configure bridge domain. */

RP/0/RSP0/CPU0:router(config)# l2vpn
RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group EVPN_SH
RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain EVPN_2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface TenGigE0/0/0/1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-ac)# exit
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface BundleEther 20.2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-ac)# storm-control broadcast pps 10000 ← Enabling
storm-control is optional
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-ac)# exit
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# evi 2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-evi)# commit

/* Configure address family session in BGP. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 200
RP/0/RSP0/CPU0:router(config-bgp)# bgp router-id 209.165.200.227
RP/0/RSP0/CPU0:router(config-bgp)# address-family l2vpn evpn

RP/0/RSP0/CPU0:router(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router(config-bgp-nbr)# description MPLSFACINGPEER
RP/0/RSP0/CPU0:router(config-bgp-nbr)# update-source Loopback 0
RP/0/RSP0/CPU0:router(config-bgp-nbr)# address-family l2vpn evpn

```

Supported Modes for EVPN Software MAC Learning


The following are the modes in which EVPN Software MAC Learning is supported:

- Single Home Device (SHD) or Single Home Network (SHN)
- Dual Home Device (DHD)—All Active Load Balancing

Single Home Device or Single Home Network Mode

The following section describes how you can configure EVPN Software MAC Learning feature in single home device or single home network (SHD/SHN) mode:

Figure 15: Single Home Device or Single Home Network Mode



In the above figure, the PE (PE1) is attached to Ethernet Segment using bundle or physical interfaces. Null Ethernet Segment Identifier (ESI) is used for SHD/SHN.

Configure EVPN in Single Home Device or Single Home Network Mode

This section describes how you can configure EVPN Software MAC Learning feature in single home device or single home network mode.

```
/* Configure bridge domain. */

RP/0/RSP0/CPU0:router(config)# l2vpn
RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group EVPN_ALL_ACTIVE
RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain EVPN_2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface Bundle-Ether1.2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# evi 2001

/* Configure advertisement of MAC routes. */

RP/0/RSP0/CPU0:router(config)# evpn
RP/0/RSP0/CPU0:router(config-evpn)# evi 2001
RP/0/RSP0/CPU0:router(config-evpn-evi)# advertise-mac

/* Configure address family session in BGP. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router#(config)# router bgp 200
RP/0/RSP0/CPU0:router#(config-bgp)# bgp router-id 09.165.200.227
RP/0/RSP0/CPU0:router#(config-bgp)# address-family l2vpn evpn
RP/0/RSP0/CPU0:router#(config-bgp)# neighbor 10.10.10.10
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# description MPLSFACING-PEER
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# update-source Loopback 0
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# address-family l2vpn evpn
```

Running Configuration

```
l2vpn
bridge group EVPN_ALL_ACTIVE
bridge-domain EVPN_2001
interface BundleEther1.2001
evi 2001
```



```

!
evpn
 evi 2001
  advertise-mac
!
router bgp 200 bgp
 router-id 40.40.40.40
 address-family l2vpn evpn
 neighbor 10.10.10.10
  remote-as 200 description MPLS-FACING-PEER
 updatesource Loopback0
 addressfamily l2vpn evpn

```

Verification

Verify EVPN in single home devices.

```
RP/0/RSP0/CPU0:router# show evpn ethernet-segment interface Te0/4/0/10 detail
```

```


Ethernet Segment Id      Interface      Nexthops
-----
N/A                      Te0/4/0/10   20.20.20.20
.....
Topology :
Operational : SH
Configured : Single-active (AApS) (default)

```

Dual Home Device—All-Active Load Balancing Mode

The following section describes how you can configure EVPN Software MAC Learning feature in dual home device (DHD) in all-active load balancing mode:

Figure 16: Dual Home Device —All-Active Load Balancing Mode



All-active load-balancing is known as Active/Active per Flow (AApF). In the above figure, identical Ethernet Segment Identifier is used on both EVPN PEs. PEs are attached to Ethernet Segment using bundle interfaces. In the CE, single bundles are configured towards two EVPN PEs. In this mode, the MAC address that is learnt is stored on both PE1 and PE2. Both PE1 and PE2 can forward the traffic within the same EVI.

Configure EVPN Software MAC Learning in Dual Home Device—All-Active Mode

This section describes how you can configure EVPN Software MAC Learning feature in dual home device—all-active mode:

```

/* Configure bridge domain. */

RP/0/RSP0/CPU0:router(config)# l2vpn
RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group EVPN_ALL_ACTIVE
RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain EVPN_2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface Bundle-Ether1.2001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# evi 2001

/* Configure advertisement of MAC routes. */

RP/0/RSP0/CPU0:router(config)# evpn
RP/0/RSP0/CPU0:router(config-evpn)# evi 2001
RP/0/RSP0/CPU0:router(config-evpn-evi)# advertise-mac
RP/0/RSP0/CPU0:router(config-evpn-evi)# exit
RP/0/RSP0/CPU0:router(config-evpn)# interface bundle-ether1
RP/0/RSP0/CPU0:router(config-evpn-ac)# ethernet-segment
RP/0/RSP0/CPU0:router(config-evpn-ac-es)# identifier type 0 01.11.00.00.00.00.00.01

/* Configure address family session in BGP. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router#(config)# router bgp 200
RP/0/RSP0/CPU0:router#(config-bgp)# bgp router-id 209.165.200.227
RP/0/RSP0/CPU0:router#(config-bgp)# address-family l2vpn evpn
RP/0/RSP0/CPU0:router#(config-bgp)# neighbor 10.10.10.10
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# remote-as 200
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# description MPLSFACING-PEER
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# update-source Loopback 0
RP/0/RSP0/CPU0:router#(config-bgp-nbr)# address-family l2vpn evpn

/* Configure Link Aggregation Control Protocol (LACP) bundle. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether1.300
RP/0/RSP0/CPU0:router(config-if)# lacp switchover suppress-flaps 300
RP/0/RSP0/CPU0:router(config-if)# exit

/* Configure VLAN Header Rewrite.*/

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface bundle-Ether1.2001 l2transport
RP/0/RSP0/CPU0:router(config-if)# encapsulation dot1q 10
RP/0/RSP0/CPU0:router(config-if)# rewrite ingress tag pop 1 symmetric

```

Running Configuration

```

l2vpn
bridge group EVPN_ALL_ACTIVE
  bridge-domain EVPN_2001
  interface Bundle-Ether1.2001
  !
  evi 2001
  !
  !

```

```

evpn
 evi 2001
 !
 advertise-mac
 !
 interface bundle-ether1
  ethernet-segment
  identifier type 0 01.11.00.00.00.00.00.01
 !
 !
router bgp 200
 bgp router-id 209.165.200.227
 address-family l2vpn evpn
 !
 neighbor 10.10.10.10
  remote-as 200
  description MPLS-FACING-PEER
  update-source Loopback0
  address-family l2vpn evpn
 !
 interface Bundle-Ether1
  lACP switchover suppress-flaps 300
  load-interval 30
 !
 interface bundle-Ether1.2001 l2transport
  encapsulation dot1q 2001
  rewrite ingress tag pop 1 symmetric
 !

```

Verification

Verify EVPN in dual home devices in All-Active mode.



Note With the EVPN IRB, the supported label mode is per-VRF.

```

RP/0/RSP0/CPU0:router# show evpn ethernet-segment interface bundle-Ether 1 carvin$

Ethernet Segment Id      Interface  Nexthops
-----
0100.211b.fce5.df00.0b00  BE11      10.10.10.10
209.165.201.1
Topology :
Operational : MHN
Configured : All-active (AAPF) (default)
Primary Services : Auto-selection
Secondary Services: Auto-selection
Service Carving Results:
Forwarders : 4003
Elected : 2002
EVI E : 2000, 2002, 36002, 36004, 36006, 36008
.....
Not Elected : 2001
EVI NE : 2001, 36001, 36003, 36005, 36007, 36009

MAC Flushing mode : Invalid

Peering timer : 3 sec [not running]
Recovery timer : 30 sec [not running]
Local SHG label : 34251

```

```
Remote SHG labels : 1
38216 : nexthop 209.165.201.1
```

Verify EVPN Software MAC Learning

Verify the packet drop statistics.



Note Disable CW configuration if any in EVPN peer nodes, as CW is not supported in EVPN Bridging.

```
RP/0/RSP0/CPU0:router# show l2vpn bridge-domain bd-name EVPN_2001 details
```

```
Bridge group: EVPN_ALL ACTIVE, bridge-domain: EVPN_2001, id: 1110,
state: up, ShgId: 0, MSTi: 0
List of EVPNs:
EVPN, state: up
evi: 2001
XC ID 0x80000458
Statistics:
packets: received 28907734874 (unicast 9697466652), sent
76882059953
bytes: received 5550285095808 (unicast 1861913597184), sent
14799781851396
MAC move: 0
List of ACs:
AC: TenGigE0/0/0/1, state is up
Type VLAN; Num Ranges: 1
...
Statistics:
packets: received 0 (multicast 0, broadcast 0, unknown
unicast 0, unicast 0), sent 45573594908
bytes: received 0 (multicast 0, broadcast 0, unknown unicast
0, unicast 0), sent 8750130222336
MAC move: 0
.....
```

Verify the EVPN EVI information with the VPN-ID and MAC address filter.

```
RP/0/RSP0/CPU0:router# show evpn evi vpn-id 2001 neighbor
```

```
Neighbor IP      vpn-id
-----
209.165.200.225  2001
209.165.201.30   2001
```

Verify the BGP L2VPN EVPN summary.

```
RP/0/RSP0/CPU0:router# show bgp l2vpn evpn summary
```

```
...
Neighbor          Spk   AS      MsgRcvd  MsgSent  TblVer   InQ   OutQ   Up/Down   St/PfxRcd
209.165.200.225   0     200     216739  229871   200781341  0     0       3d00h    348032
209.165.201.30    0     200     6462962 4208831  200781341 10     0       2d22h    35750
```

Verify the MAC updates to the L2FIB table in a line card.

```
RP/0/RSP0/CPU0:router# show l2vpn mac mac all location
```

```

Topo ID Producer Next Hop(s)      Mac Address      IP Address
-----
1112     0/6/CPU0 Te0/0/0/1 00a3.0001.0001

```

Verify the MAC updates to the L2FIB table in a route switch processor (RSP).

```
RP/0/RSP0/CPU0:router# show l2vpn mac mac all location
```

```

Topo ID  Producer Next Hop(s)      Mac Address      IP Address
-----
1112     0/6/CPU0 0/0/0/1 00a3.0001.0001

```

Verify the summary information for the MAC address.

```
RP/0/RSP0/CPU0:router# show l2vpn forwarding bridge-domain EVPN_ALL_ACTIVE:EVPN_2001
mac-address location
```

```

.....
Mac Address      Type          Learned from/Filtered on  LC learned  Resync Age/Last Change
Mapped to
0000.2001.5555  dynamic      Te0/0/0/2                N/A         11 Jan 14:37:22
N/A <-- local dynamic
00bb.2001.0001  dynamic      Te0/0/0/2                N/A         11 Jan 14:37:22
N/A
0000.2001.1111  EVPN         BD id: 1110              N/A         N/A
N/A <-- remote static
00a9.2002.0001  EVPN         BD id: 1110              N/A         N/A
N/A

```

Verify the EVPN EVI information with the VPN-ID and MAC address filter.

```
RP/0/RSP0/CPU0:router# show evpn evi vpn-id 2001 mac
```

```

EVI      MAC address      IP address      Nexthop      Label
----      -
2001     00a9.2002.0001  ::              10.10.10.10  34226        <-- Remote MAC
2001     00a9.2002.0001  ::              209.165.201.30 34202
2001     0000.2001.5555  20.1.5.55      TenGigE0/0/0/2 34203        <-- local MAC

```

```
RP/0/RSP0/CPU0:router# RP/0/RSP0/CPU0:router# show evpn evi vpn-id 2001 mac 00a9.2002.0001
detail
```

```

EVI      MAC address      IP address      Nexthop      Label
----      -
2001     00a9.2002.0001  ::              10.10.10.10  34226
2001     00a9.2002.0001  ::              209.165.201.30 34202

```

```

Ethernet Tag : 0
Multi-paths Resolved : True <--- aliasing to two remote PE with All-Active load balancing

```

```

Static : No
Local Ethernet Segment : N/A
Remote Ethernet Segment : 0100.211b.fce5.df00.0b00
Local Sequence Number : N/A
Remote Sequence Number : 0
Local Encapsulation : N/A

```

```
Remote Encapsulation : MPLS
```

Verify the BGP routes associated with EVPN with bridge-domain filter.

```
RP/0/RSP0/CPU0:router# show bgp l2vpn evpn bridge-domain EVPN_2001 route-type 2

*> [2][0][48][00bb.2001.0001][0]/104
      0.0.0.0          0 i <----- locally learnt MAC
*>i[2][0][48][00a9.2002.00be][0]/104
      10.10.10.10    0 i <----- remotely learnt MAC
* i 209.165.201.30 100 0 i
```

EVPN Out of Service

The EVPN Out of Service feature enables you to control the state of bundle interfaces that are part of an Ethernet segment that have Link Aggregation Control protocol (LACP) configured. This feature enables you to put a node out of service (OOS) without having to manually shutdown all the bundles on their provider edge (PE).


Use the **cost-out** command to bring down all the bundle interfaces belonging to an Ethernet VPN (EVPN) Ethernet segment on a node. The Ethernet A-D Ethernet Segment (ES-EAD) routes are withdrawn before shutting down the bundles. The PE signals to the connected customer edge (CE) device to bring down the corresponding bundle member. This steers away traffic from this PE node without traffic disruption. The traffic that is bound for the Ethernet segment from the CE is directed to the peer PE in a multi-homing environment.



Note EVPN cost-out is supported only on manually configured ESIs.

In the following topology, the CE is connected to PE1 and PE2. When you configure the **cost-out** command on PE1, all the bundle interfaces on the Ethernet segment are brought down. Also, the corresponding bundle member is brought down on the CE. Hence, the traffic for this Ethernet segment is now sent to PE2 from the CE.

Figure 17: EVPN Out of Service



To bring up the node into service, use **no cost-out** command. This brings up all the bundle interfaces belonging to EVPN Ethernet segment on the PE and the corresponding bundle members on the CE.

When the node is in cost-out state, adding a new bundle Ethernet segment brings that bundle down. Similarly, removing the bundle Ethernet segment brings that bundle up.

Use **startup-cost-in** command to bring up the node into service after the specified time on reload. The node will cost-out when EVPN is initialized and remain cost-out until the set time. If you execute **evpn no startup-cost-in** command while timer is running, the timer stops and node is cost-in.

The 'cost-out' configuration always takes precedence over the 'startup-cost-in' timer. So, if you reload with both the configurations, cost-out state is controlled by the 'cost-out' configuration and the timer is not relevant. Similarly, if you reload with the startup timer, and configure 'cost-out' while timer is running, the timer is stopped and OOS state is controlled only by the 'cost-out' configuration.

If you do a proc restart while the startup-cost-in timer is running, the node remains in cost-out state and the timer restarts.

Configure EVPN Out of Service

This section describes how you can configure EVPN Out of Service.

```
/* Configuring node cost-out on a PE */

Router# configure
Router(config)# evpn
Router(config-evpn)# cost-out
Router(config-evpn) commit

/* Bringing up the node into service */

Router# configure
Router(config)# evpn
Router(config-evpn)# no cost-out
Router(config-evpn) commit

/* Configuring the timer to bring up the node into service after the specified time on
reload */

Router# configure
Router(config)# evpn
Router(config-evpn)# startup-cost-in 6000
Router(config-evpn) commit
```

Running Configuration

```
configure
evpn
  cost-out
!

configure
evpn
  startup-cost-in 6000
!
```

Verification

Verify the EVPN Out of Service configuration.

```

/* Verify the node cost-out configuration */

Router# show evpn summary
Fri Apr 7 07:45:22.311 IST
Global Information
-----
Number of EVIs : 2
Number of Local EAD Entries : 0
Number of Remote EAD Entries : 0
Number of Local MAC Routes : 0
Number of Local MAC Routes : 5
      MAC : 5
      MAC-IPv4 : 0
      MAC-IPv6 : 0
Number of Local ES:Global MAC : 12
Number of Remote MAC Routes : 7
      MAC : 7
      MAC-IPv4 : 0
      MAC-IPv6 : 0
Number of Local IMCAST Routes : 56
Number of Remote IMCAST Routes: 56
Number of Internal Labels : 5
Number of ES Entries : 9
Number of Neighbor Entries : 1
EVPN Router ID : 192.168.0.1
BGP Router ID : ::
BGP ASN : 100
PBB BSA MAC address : 0207.1fee.be00
Global peering timer : 3 seconds
Global recovery timer : 30 seconds
EVPN cost-out : TRUE
      startup-cost-in timer : Not configured

```

```

/* Verify the no cost-out configuration */

Router# show evpn summary
Fri Apr 7 07:45:22.311 IST
Global Information
-----
Number of EVIs : 2
Number of Local EAD Entries : 0
Number of Remote EAD Entries : 0
Number of Local MAC Routes : 0
Number of Local MAC Routes : 5
      MAC : 5
      MAC-IPv4 : 0
      MAC-IPv6 : 0
Number of Local ES:Global MAC : 12
Number of Remote MAC Routes : 7
      MAC : 7
      MAC-IPv4 : 0
      MAC-IPv6 : 0
Number of Local IMCAST Routes : 56
Number of Remote IMCAST Routes: 56
Number of Internal Labels : 5
Number of ES Entries : 9
Number of Neighbor Entries : 1
EVPN Router ID : 192.168.0.1
BGP Router ID : ::
BGP ASN : 100
PBB BSA MAC address : 0207.1fee.be00
Global peering timer : 3 seconds

```



```

Global recovery timer      :      30 seconds
EVPN cost-out             : FALSE
    startup-cost-in timer : Not configured

/* Verify the startup-cost-in timer configuration */

Router# show evpn summary
Fri Apr 7 07:45:22.311 IST
Global Information
-----
Number of EVIs           : 2
Number of Local EAD Entries : 0
Number of Remote EAD Entries : 0
Number of Local MAC Routes : 0
Number of Local MAC Routes : 5
    MAC                   : 5
    MAC-IPv4               : 0
    MAC-IPv6               : 0
Number of Local ES:Global MAC : 12
Number of Remote MAC Routes : 7
    MAC                   : 7
    MAC-IPv4               : 0
    MAC-IPv6               : 0
Number of Local IMCAST Routes : 56
Number of Remote IMCAST Routes: 56
Number of Internal Labels   : 5
Number of ES Entries        : 9
Number of Neighbor Entries  : 1
EVPN Router ID              : 192.168.0.1
BGP Router ID               : ::
BGP ASN                     : 100
PBB BSA MAC address         : 0207.1fee.be00
Global peering timer        :      3 seconds
Global recovery timer       :      30 seconds
EVPN node cost-out          : TRUE
    startup-cost-in timer   : 6000

```

EVPN Multiple Services per Ethernet Segment

EVPN Multiple Services per Ethernet Segment feature allows you to configure multiple services over single Ethernet Segment (ES). Instead of configuring multiple services over multiple ES, you can configure multiple services over a single ES.

You can configure the following services on a single Ethernet Bundle; you can configure one service on each sub-interface.

- Flexible cross-connect (FXC) service. It supports VLAN Unaware, VLAN Aware, and Local Switching modes.

For more information, see *Configure Point-to-Point Layer 2 Services* chapter in *L2VPN and Ethernet Services Configuration Guide for Cisco NCS 5500540 Series Routers*.

- EVPN-VPWS Xconnect service

For more information, see *EVPN Virtual Private Wire Service (VPWS)* chapter in *L2VPN and Ethernet Services Configuration Guide for Cisco NCS 5500540 Series Routers*.

- EVPN Integrated Routing and Bridging (IRB)

For more information, see *Configure EVPN IRB* chapter in *L2VPN and Ethernet Services Configuration Guide for Cisco NCS 5500540 Series Routers*.

- Native EVPN

For more information see, *EVPN Features* chapter in *L2VPN and Ethernet Services Configuration Guide for Cisco NCS 5500540 Series Routers*.

All these services are supported only on all-active multihoming scenario.

Configure EVPN Multiple Services per Ethernet Segment

Consider a customer edge (CE) device connected to two provider edge (PE) devices through Ethernet Bundle interface 22001. Configure multiple services on Bundle Ethernet sub-interfaces.

Configuration Example

Consider Bundle-Ether22001 ES, and configure multiple services on sub-interface.

```
/* Configure attachment circuits */
Router# configure
Router(config)# interface Bundle-Ether22001.12 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 12
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether22001.13 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 13
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether22001.14 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 14
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether22001.1 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 1
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether22001.2 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 2
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether22001.3 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 3
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit
Router(config)# interface Bundle-Ether22001.4 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 4
Router(config-l2vpn-subif)# exit
Router(config-l2vpn)# exit

/*Configure VLAN Unaware FXC Service */
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware fxc_mh1
Router(config-l2vpn-fxs-vu)# interface Bundle-Ether22001.1
Router(config-l2vpn-fxs-vu)# interface Bundle-Ether22001.2
Router(config-l2vpn-fxs-vu)# interface Bundle-Ether22001.3
Router(config-l2vpn-fxs-vu)# neighbor evpn evi 21006 target 22016
Router(config-l2vpn-fxs-vu)# commit
```

```

/* Configure VLAN Aware FXC Service */
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 24001
Router(config-l2vpn-fxs-va)# interface Bundle-Ether22001.12
Router(config-l2vpn-fxs-va)# interface Bundle-Ether22001.13
Router(config-l2vpn-fxs-va)# interface Bundle-Ether22001.14
Router(config-l2vpn-fxs-va)# commit

/* Configure Local Switching - Local switching is supported only on VLAN-aware FXC */
PE1
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 31400
Router(config-l2vpn-fxs-va)# interface Bundle-Ether22001.1400
Router(config-l2vpn-fxs-va)# interface Bundle-Ether23001.1400
Router(config-l2vpn-fxs-va)# commit
Router(config-l2vpn-fxs)# exit
PE2
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-aware evi 31401
Router(config-l2vpn-fxs-va)# interface Bundle-Ether22001.1401
Router(config-l2vpn-fxs-va)# interface Bundle-Ether23001.1401
Router(config-l2vpn-fxs-va)# commit
Router(config-l2vpn-fxs)# exit

/* Configure EVPN-VPWS xconnect service and native EVPN with IRB */

Router# configure
Router(config)# interface Bundle-Ether22001.11 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 11
Router(config-l2vpn-subif)# rewrite ingress tag pop 2 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit

Router# configure
Router(config)# interface Bundle-Ether22001.21 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 1 second-dot1q 21
Router(config-l2vpn-subif)# rewrite ingress tag pop 2 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit

Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group xg22001
Router(config-l2vpn-xc)# p2p evpn-vpws-mclag-22001
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether22001.11
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 22101 target 220101 source 220301
Router(config-l2vpn-xc-p2p-pw)# commit
Router(config-l2vpn-xc-p2p-pw)# exit

Router # configure
Router (config)# l2vpn
Router (config-l2vpn)# bridge group native_evpn1
Router (config-l2vpn-bg)# bridge-domain bd21
Router (config-l2vpn-bg-bd)# interface Bundle-Ether22001.21
Router (config-l2vpn-bg-bd-ac)# routed interface BVI21
Router (config-l2vpn-bg-bd-bvi)# evi 22021
Router (config-l2vpn-bg-bd-bvi)# commit
Router (config-l2vpn-bg-bd-bvi)# exit

/* Configure Native EVPN */

```

```

Router # configure
Router (config)# evpn
Router (config-evpn)# interface Bundle-Ether22001
Router (config-evpn-ac)# ethernet-segment identifier type 0 ff.ff.ff.ff.ff.ff.ff.ee
Router (config-evpn-ac-es)# bgp route-target 2200.0001.0001
Router (config-evpn-ac-es)# exit
Router (config-evpn)# evi 24001
Router (config-evpn-evi)# bgp
Router (config-evpn-evi-bgp)# route-target import 64:24001
Router (config-evpn-evi-bgp)# route-target export 64:24001
Router (config-evpn-evi-bgp)# exit
Router (config-evpn-evi)# exit
Router (config-evpn)# evi 21006
Router (config-evpn-evi)# bgp
Router (config-evpn-evi-bgp)# route-target route-target 64:10000
Router (config-evpn-evi-bgp)# exit
Router (config-evpn-evi)# exit
Router (config-evpn)# evi 22101
Router (config-evpn-evi)# bgp
Router (config-evpn-evi-bgp)# route-target import 64:22101
Router (config-evpn-evi-bgp)# route-target export 64:22101
Router (config-evpn-evi-bgp)# exit
Router (config-evpn-evi)# exit
Router (config-evpn)# evi 22021
Router (config-evpn-evi)# bgp
Router (config-evpn-evi-bgp)# route-target import 64: 22021
Router (config-evpn-evi-bgp)# route-target export 64: 22021
Router (config-evpn-evi-bgp)# exit
Router (config-evpn-evi)# exit
Router (config-evpn-evi)# advertise-mac
Router (config-evpn-evi)# exit
Router (config-evpn)# evi 22022
Router (config-evpn-evi)# bgp
Router (config-evpn-evi-bgp)# route-target import 64: 22022
Router (config-evpn-evi-bgp)# route-target export 64: 22022
Router (config-evpn-evi-bgp)# exit
Router (config-evpn-evi)# advertise-mac
Router (config-evpn-evi)# commit
Router (config-evpn-evi)# exit

```

Running Configuration

```

/* Configure attachment circuits */
interface Bundle-Ether22001.12 l2transport
encapsulation dot1q 1 second-dot1q 12
!
interface Bundle-Ether22001.13 l2transport
encapsulation dot1q 1 second-dot1q 13
!
interface Bundle-Ether22001.14 l2transport
encapsulation dot1q 1 second-dot1q 14
!
interface Bundle-Ether22001.1 l2transport
encapsulation dot1q 1 second-dot1q 1
!
interface Bundle-Ether22001.2 l2transport
encapsulation dot1q 1 second-dot1q 2
!
interface Bundle-Ether22001.3 l2transport
encapsulation dot1q 1 second-dot1q 3

```

```

!
interface Bundle-Ether22001.4 l2transport
encapsulation dot1q 1 second-dot1q 4

/*Configure VLAN Unaware FXC Service */
flexible-xconnect-service vlan-unaware fxc_mh1
  interface Bundle-Ether22001.1
  interface Bundle-Ether22001.2
  interface Bundle-Ether22001.3
  neighbor evpn evi 21006 target 22016
!
/*Configure VLAN Aware FXC Service */
l2vpn
flexible-xconnect-service vlan-aware evi 24001
  interface Bundle-Ether22001.12
  interface Bundle-Ether22001.13
  interface Bundle-Ether22001.14

/* Configure Local Switching */
flexible-xconnect-service vlan-aware evi 31400
  interface Bundle-Ether22001.1400
  interface Bundle-Ether23001.1400
!
flexible-xconnect-service vlan-aware evi 31401
  interface Bundle-Ether22001.1401
  interface Bundle-Ether23001.1401
!

/* Configure EVPN-VPWS xconnect service and native EVPN with IRB */
interface Bundle-Ether22001.11 l2transport
  encapsulation dot1q 1 second-dot1q 11
  rewrite ingress tag pop 2 symmetric
!
interface Bundle-Ether22001.21 l2transport
  encapsulation dot1q 1 second-dot1q 21
  rewrite ingress tag pop 2 symmetric
!
!
l2vpn
xconnect group xg22001
p2p evpn-vpws-mclag-22001
  interface Bundle-Ether22001.11
  neighbor evpn evi 22101 target 220101 source 220301
!
bridge group native_evpn1
  bridge-domain bd21
  interface Bundle-Ether22001.21
  routed interface BVI21
  evi 22021
!
/* Configure Native EVPN */
Evpn
interface Bundle-Ether22001
  ethernet-segment identifier type 0 ff.ff.ff.ff.ff.ff.ff.ff.00
  bgp route-target 2200.0001.0001
!
  evi 24001
  bgp
    route-target import 64:24001
    route-target export 64:24001
!
  evi 21006
  bgp
    route-target 64:100006

```

```

!
evi 22101
  bgp
    route-target import 64:22101
    route-target export 64:22101
!
evi 22021
  bgp
    route-target import 64:22021
    route-target export 64:22021
!
  advertise-mac
!
evi 22022
  bgp
    route-target import 64:22022
    route-target export 64:22022
!
  advertise-mac
!

```

Verification

Verify if each of the services is configured on the sub-interface.

```

Router# show l2vpn xconnect summary
Number of groups: 6
Number of xconnects: 505 Up: 505 Down: 0 Unresolved: 0 Partially-programmed: 0
AC-PW: 505 AC-AC: 0 PW-PW: 0 Monitor-Session-PW: 0
Number of Admin Down segments: 0
Number of MP2MP xconnects: 0
  Up 0 Down 0
Advertised: 0 Non-Advertised: 0

```

```

Router# show l2vpn flexible-xconnect-service summary
Number of flexible xconnect services: 74
Up: 74

```

```

Router# show l2vpn flexible-xconnect-service name fxc_mh1
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
SB = Standby, SR = Standby Ready, (PP) = Partially Programmed
Flexible XConnect Service Segment
Name      ST  Type  Description  ST
-----
fxc_mh1  UP  AC:   BE22001.1   UP
          AC:   BE22001.2   UP
          AC:   BE22001.3   UP
-----

```

```

Router# show l2vpn flexible-xconnect-service name evi:24001

```

```

Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
SB = Standby, SR = Standby Ready, (PP) = Partially Programmed
Flexible XConnect Service Segment
Name      ST  Type  Description  ST
-----
evi:24001 UP  AC:   BE22001.11  UP
          AC:   BE22001.12  UP
          AC:   BE22001.13  UP
-----

```

```
AC: BE22001.14 UP
```

```
-----
Router# show l2vpn xconnect group xg22001 xc-name evpn-vpws-mclag-22001
Fri Sep 1 17:28:58.259 UTC
Legend: ST = State, UP = Up, DN = Down, AD = Admin Down, UR = Unresolved,
SB = Standby, SR = Standby Ready, (PP) = Partially Programmed
XConnect
Group      Name                               ST      Description ST      Segment 2
          Name                               ST      Description ST      Description                               ST
-----
xg22001   evpn-vpws-mclag-22001             UP      BE22001.101 UP      EVPN 22101, 220101, 64.1.1.6 UP
-----
```

Associated Commands

- evpn
- evi
- ethernet-segment
- advertise-mac
- show evpn ethernet-segment
- show evpn evi
- show evpn summary
- show l2vpn xconnect summary
- show l2vpn flexible-xconnect-service
- show l2vpn xconnect group

Network Convergence using Core Isolation Protection

The Network Convergence using Core Isolation Protection feature allows the router to converge fast when remote links and local interfaces fail. This feature reduces the duration of traffic drop by rapidly rerouting traffic to alternate paths. This feature uses Object Tracking (OT) to detect remote link failure and failure of connected interfaces.

Tracking interfaces can only detect failure of connected interfaces and not failure of a remote router interfaces that provides connectivity to the core. Tracking one or more BGP neighbor sessions along with one or more of the neighbor's address-families enables you to detect remote link failure.

Object Tracking

Object tracking (OT) is a mechanism for tracking an object to take any client action on another object as configured by the client. The object on which the client action is performed may not have any relationship to the tracked objects. The client actions are performed based on changes to the properties of the object being tracked.

You can identify each tracked object by a unique name that is specified by the track command in the configuration mode.

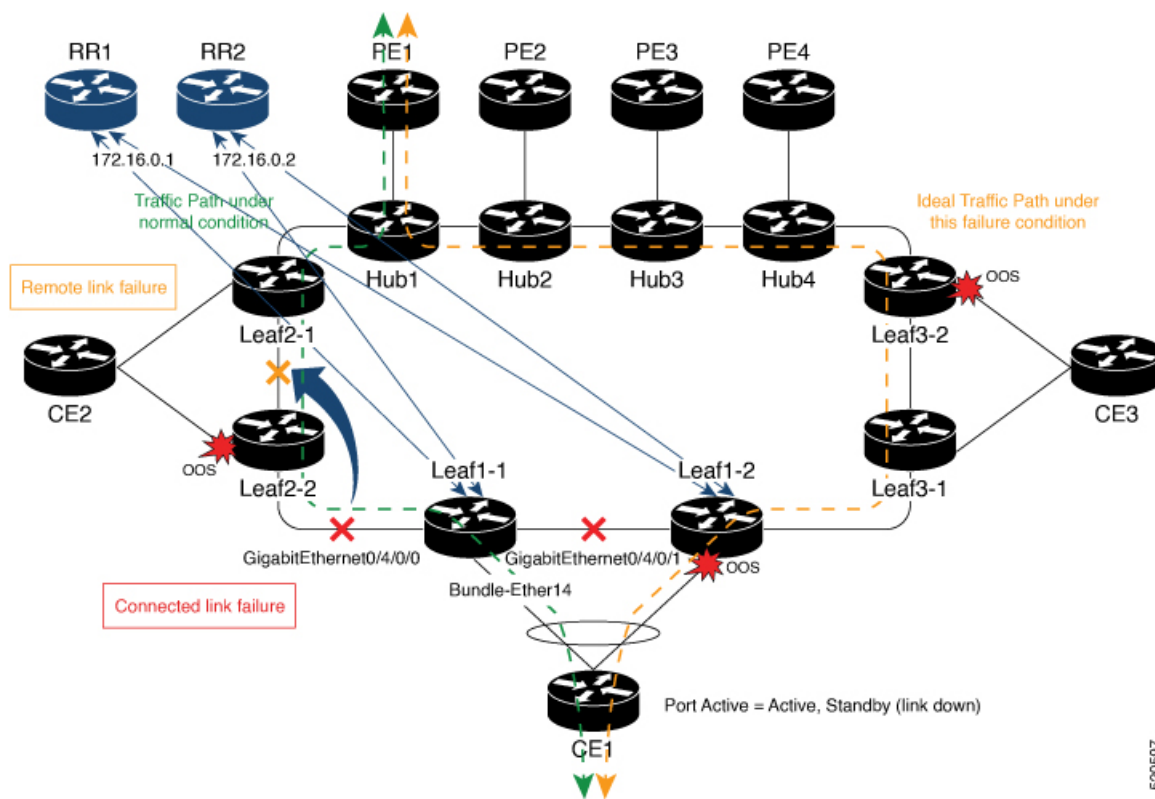
The tracking process receives the notification when the tracked object changes its state. The state of the tracked objects can be up or down.

You can also track multiple objects by a list. You can use a flexible method for combining objects with Boolean logic. This functionality includes:

- Boolean AND function—When a tracked list has been assigned a Boolean AND function, each object defined within a subset must be in an up state, so that the tracked object can also be in the up state.
- Boolean OR function—When the tracked list has been assigned a Boolean OR function, it means that at least one object defined within a subset must also be in an up state, so that the tracked object can also be in the up state.

For more information on OT, see the *Configuring Object Tracking* chapter in the *System Management Configuration Guide for Cisco NCS 5500 Series Routers*.

Figure 18: EVPN Convergence Using Core Isolation Protection



Consider a traffic flow from CE1 to PE1. The CE1 can send the traffic either from Leaf1-1 or Leaf1-2. When Leaf1-1 loses the connectivity to both the local links and remote link, BGP sessions to both route reflectors (RRs) are down; the Leaf1-1 brings down the Bundle-Ether14 connected to CE1. The CE1 redirects the traffic from Leaf1-2 to PE1.

You can track the connected interfaces to identify the connected link failures. However, if there is a remote link failure, tracking connected interfaces does not identify the remote link failures. You must track BGP sessions to identify the remote link failure.



Note When you configure the **bgp graceful-restart** command, unconfiguring a neighbor is considered as a non-gr event. This generates a BGP notification to the neighbor before the neighbor is unconfigured.

On the remote router, if the track is configured for this neighbor, the track state is brought down immediately.

However, certain configurations are treated as graceful reset reason and when unconfigured they suppress the BGP notification to the neighbor. The route-reflector-client configuration under the neighbor or neighbor address-family is one of the examples.

On the remote router, if the track is configured for this neighbor, the track state is not brought down immediately because a notification is not received.

To overcome this situation, shutdown the neighbor before unconfiguring the neighbor. This generates a BGP notification to the neighbor, and any track configured for the neighbor is brought down immediately.

Configure EVPN Convergence using Core Isolation Protection

A tracked list contains one or more objects. The Boolean expression enables tracking objects using either AND or OR operators. For example, when tracking two interfaces, using the AND operator, up means that *both* interfaces are up, and down means that *either* interface is down.



Note An object must exist before it can be added to a tracked list.

The NOT operator is specified for one or more objects and negates the state of the object.

After configuring the tracked object, you must associate the neighbor or interface whose state must be tracked.

Perform the following tasks to configure EVPN convergence using core isolation protection:

- Configure BGP
- Track the Line Protocol State of an Interface
- Track neighbor address-family state
- Track objects for both interfaces and neighbors

Configuration Example

In this example, Leaf1-1 brings the down the AC connected to CE1 when:

Both local interfaces GigabitEthernet0/4/0/0 and GigabitEthernet0/4/0/1 of Leaf1-1 are down.

OR

Leaf1-1 BGP sessions to both RRs are down.

CE1 re-directs the traffic it was sending to Leaf1-1 to Leaf1-2.

Perform the following tasks on Leaf1-1:

```
/* Configure BGP */
Router# configure
```

```

Router(config)# router bgp 100
Router(config-bgp)# address-family l2vpn evpn
Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 172.16.0.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr-af)# neighbor 172.16.0.1
Router(config-bgp-nbr)# remote-as 100
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr-af)# commit

/* Track the Line Protocol State of an Interface */
Router# configure
Router(config)# track interface-1
Router(config-track)# type line-protocol state
Router(config-track-line-prot)# interface GigabitEthernet0/4/0/0
Router(config-track-line-prot)#exit
Router(config-track)#exit
Router(config)# track interface-2
Router(config-track)# type line-protocol state
Router(config-track-line-prot)# interface GigabitEthernet0/4/0/1
Router(config-track-line-prot)#exit
Router(config-track)#exit
Router(config)# track interface-group-1
Router(config-track)# type list boolean or
Router(config-track-list-boolean)# object interface-1
Router(config-track-list-boolean)# object interface-2
Router(config-track-list-boolean)# commit

/* Track neighbor address-family state */
Router# configure
Router(config)# track neighbor-A
Router(config-track)# type bgp neighbor address-family state
Router(config-track-bgp-nbr-af)# address-family l2vpn evpn
Router(config-track-bgp-neighbor)# neighbor 172.16.0.1
Router(config-track-bgp-neighbor)# exit
Router(config-track-bgp-nbr-af)# exit
Router(config-track)# exit
Router(config)# track neighbor-B
Router(config-track)# type bgp neighbor address-family state
Router(config-track-bgp-nbr-af)# address-family l2vpn evpn
Router(config-track-bgp-neighbor)# neighbor 172.16.0.2
Router(config-track-bgp-neighbor)# exit
Router(config-track-bgp-nbr-af)# exit
Router(config-track)# exit
Router(config)# track neighbor-group-1
Router(config-track)# type list boolean or
Router(config-track-list-boolean)# object neighbor-A
Router(config-track-list-boolean)# object neighbor-B
Router(config-track-list-boolean)# commit

/* Track objects for both interfaces and neighbors */
Router# configure
Router(config)# track core-group-1
Router(config-track)# type list boolean and
Router(config-track-list-boolean)# object neighbor-group-1
Router(config-track-list-boolean)# object interface-group-1
Router(config-track-list-boolean)# action
Router(config-track-action)# track-down error-disable interface Bundle-Ether14 auto-recover
Router(config-track-action)# commit

```

Running Configuration

This section shows EVPN convergence using core isolation protection running configuration.

```
router bgp 100
  address-family l2vpn evpn
  !
  neighbor 172.16.0.1
    remote-as 100
    address-family l2vpn evpn
  !
  !
  neighbor 172.16.0.2
    remote-as 100
    address-family l2vpn evpn
  !
  !
!

track interface-1
  type line-protocol state
  interface GigabitEthernet0/4/0/0
  !
!

track interface-2
  type line-protocol state
  interface GigabitEthernet0/4/0/1
  !
!

track interface-group-1
  type list boolean or
  object interface-1
  object interface-2
  !
!

track neighbor-A
  type bgp neighbor address-family state
  address-family l2vpn evpn
  neighbor 172.16.0.1
  !
!

track neighbor-B
  type bgp neighbor address-family state
  address-family l2vpn evpn
  neighbor 172.16.0.1
  !
!

track neighbor-group-1
  type list boolean or
  object neighbor-A
  object neighbor-B
  !
!

track core-group-1
  type list boolean and
  object neighbor-group-1
  object interface-group-1
  !
action
```

```

track-down error-disable interface Bundle-Ether14 auto-recover
!
!

```

Verification

Verify that you have configured the EVPN convergence using core isolation protection feature successfully.

```
Router# show track
```

```
Wed May 27 04:42:11.995 UTC
```

```
Track neighbor-A
```

```

BGP Neighbor AF L2VPN EVPN NBR 172.16.0.1 vrf default
Reachability is UP
  Neighbor Address Reachability is Up
  BGP Neighbor Address-family state is Up
4 changes, last change UTC Tue May 26 2020 20:14:33.171

```

```
Track neighbor-B
```

```

BGP Neighbor AF L2VPN EVPN NBR 172.16.0.2 vrf default
Reachability is UP
  Neighbor Address Reachability is Up
  BGP Neighbor Address-family state is Up
4 changes, last change UTC Tue May 26 2020 20:14:27.527

```

```
Track core-group-1
```

```

List boolean and is UP
2 changes, last change 20:14:27 UTC Tue May 26 2020
  object interface-group-1 UP
  object neighbor-group-1 UP

```

```
Track interface-1
```

```

Interface GigabitEthernet0/4/0/0 line-protocol
Line protocol is UP
2 changes, last change 20:13:32 UTC Tue May 26 2020

```

```
Track interface-2
```

```

Interface GigabitEthernet0/4/0/1 line-protocol
Line protocol is UP
2 changes, last change 20:13:28 UTC Tue May 26 2020

```

```
Track interface-group-1
```

```

List boolean or is UP
2 changes, last change 20:13:28 UTC Tue May 26 2020
  object interface-2 UP
  object interface-1 UP

```

```
Track neighbor-group-1
```

```

List boolean or is UP
2 changes, last change 20:14:27 UTC Tue May 26 2020
  object neighbor-A UP
  object neighbor-B UP

```

```
Router# show track brief
```

```
Wed May 27 04:39:19.740 UTC
```

Track	Object	Parameter
neighbor-A Up	bgp nbr L2VPN EVPN 172.16.0.1 vrf defau	reachability
neighbor-B Up	bgp nbr L2VPN EVPN 172.16.0.1 vrf defau	reachability

```

core-group-1                list                boolean and
  Up
interface-1                 interface GigabitEthernet0/4/0/0  line protocol
  Up
interface-2                 interface GigabitEthernet0/4/0/1  line protocol
  Up
interface-group-1          list                boolean or
  Up
neighbor-group-1          list                boolean or
  Up

```

```

Router# show bgp track
Wed May 27 05:05:51.285 UTC

```

VRF	Address-family	Neighbor	Status	Flags
default	L2VPN EVPN	172.16.0.1	UP	0x01
default	L2VPN EVPN	172.16.0.2	UP	0x01

```


Processed 2 entries

```

Conditional Advertisement of Default-Originate

The router advertises the default-originate (0.0.0.0/0) towards the network fabric only upon receiving all the core routes. The router withdraws the advertisement of default-originate when the core is isolated. To avoid traffic drop, install the routes in the hardware. To accommodate an additional delay for the routes to be installed in the hardware, you can configure a timeout for the installed routes.

Figure 19: Advertisement of default-originate



In this topology, PE3 advertises the default-originate to CE only when the PE3 session to RR is established and all the routes are received from the RR.

Configure Conditional Advertisement of Default-Originate

Perform the following tasks to configure conditional advertisement of default-originate.

- Configure BGP
- Configure RPL
- Track BGP neighbor address-family state

Configuration Example

Perform the following task on PE3:

```
/* Configure BGP */
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# bgp router-id 192.0.2.1
Router(config-bgp)# address-family vpnv4 unicast
```

```

Router(config-bgp-af)# exit
Router(config-bgp)# neighbor 172.16.0.5
Router(config-bgp-nbr)# remote-as 200
Router(config-bgp-nbr)# address-family vpnv4 unicast
Router(config-bgp-nbr-af)# exit
Router(config-bgp-nbr)# exit
Router(config-bgp)# vrf cust1
Router(config-bgp-vrf)# rd auto
Router(config-bgp-vrf)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# redistribute connected
Router(config-bgp-vrf-af)# redistribute static
Router(config-bgp-vrf-af)# exit
Router(config-bgp-vrf)# neighbor 172.16.0.5
Router(config-bgp-vrf-nbr)# remote-as 200
Router(config-bgp-vrf-nbr)# address-family ipv4 unicast
Router(config-bgp-vrf-nbr-af)# default-originate route-policy track-bgp-core-policy
Router(config-bgp-vrf-nbr-af)# route-policy pass in
Router(config-bgp-vrf-nbr-af)# route-policy pass out
Router(config-bgp-vrf-nbr-af)# commit

/* Configure RPL */
Router# configure
Router(config)# route-policy track-bgp-core-policy
Router(config-rpl)# if track core-group-1 is up then
Router(config-rpl-if)# pass
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy
Router(config)# commit

/* Track BGP neighbor address-family state */
Router# configure
Router(config)# track core-group-1
Router(config-track)# type bgp neighbor address-family state
Router(config-track-bgp-nbr-af)# address-family vpnv4 unicast
Router(config-track-bgp-neighbor)# neighbor 172.16.0.5
Router(config-track-bgp-neighbor)# commit

```

Running Configuration

This section shows conditional advertisement of default-originate running configuration.

```

configure
router bgp 100
  bgp router-id 192.0.2.1
  address-family vpnv4 unicast
!
  neighbor 172.16.0.5
  remote-as 200
  address-family vpnv4 unicast
!

vrf cust1
  rd auto
  address-family ipv4 unicast
  redistribute connected
  redistribute static
!
neighbor 172.16.0.5
remote-as 200
address-family ipv4 unicast
  default-originate route-policy track-bgp-core-policy

```

```

    route-policy pass in
    route-policy pass out
!

route-policy track-bgp-core-policy
if track core-group-1 is up then
    pass
endif
end-policy
!
track network-core
type bgp neighbor address-family state
address-family vpnv4 unicast
neighbor 172.16.0.5
!

```

Verification

Verify conditional advertisement of default-originate.

```

Router# show rpl active route-policy
Wed May 27 06:54:31.902 UTC

```

```

ACTIVE -- Referenced by at least one policy which is attached
INACTIVE -- Only referenced by policies which are not attached
UNUSED -- Not attached (directly or indirectly) and not referenced

```

The following policies are (ACTIVE)

```

-----
    track-bgp-core
-----

Router# show rpl route-policy track-bgp-core-policy
Wed May 27 06:54:38.090 UTC
route-policy track-bgp-core-policy
if track core-group-1 is up then
    pass
endif
end-policy
!

```

```

Router# show bgp policy route-policy track-bgp-core-policy summary
Wed May 27 06:54:42.823 UTC
Network          Next Hop          From              Advertised to
0.0.0.0/0        0.0.0.0           Local             172.16.0.5

```

```

Router# show bgp neighbor 172.16.0.5
Wed May 27 06:55:39.535 UTC

```

```

BGP neighbor is 172.16.0.5
Remote AS 9730, local AS 9730, internal link
Remote router ID 172.16.0.5
BGP state = Established, up for 10:41:12
[snip]
For Address Family: IPv4 Unicast
BGP neighbor version 2
Update group: 0.4 Filter-group: 0.1 No Refresh request being processed
Default information originate: default route-policy track-bgp-core-policy, default sent
| AF-dependent capabilities:
[snip]
Track Enabled, Status UP, Nbr GR state Not Enabled, EOR tmr Not Running
Advertise routes with local-label via Unicast SAFI

```


EVPN Routing Policy

The EVPN Routing Policy feature provides the route policy support for address-family L2VPN EVPN. This feature adds EVPN route filtering capabilities to the routing policy language (RPL). The filtering is based on various EVPN attributes.

A routing policy instructs the router to inspect routes, filter them, and potentially modify their attributes as they are accepted from a peer, advertised to a peer, or redistributed from one routing protocol to another.

This feature enables you to configure route-policies using EVPN network layer reachability information (NLRI) attributes of EVPN route type 1 to 5 in the route-policy match criteria, which provides more granular definition of route-policy. For example, you can specify a route-policy to be applied to only certain EVPN route-types or any combination of EVPN NLRI attributes. This feature provides flexibility in configuring and deploying solutions by enabling route-policy to filter on EVPN NLRI attributes.

To implement this feature, you need to understand the following concepts:

- Routing Policy Language
- Routing Policy Language Structure
- Routing Policy Language Components
- Routing Policy Language Usage
- Policy Definitions
- Parameterization
- Semantics of Policy Application
- Policy Statements
- Attach Points

For information on these concepts, see [Implementing Routing Policy](#).

Currently, this feature is supported only on BGP neighbor "in" and "out" attach points. The route policy can be applied only on inbound or outbound on a BGP neighbor.

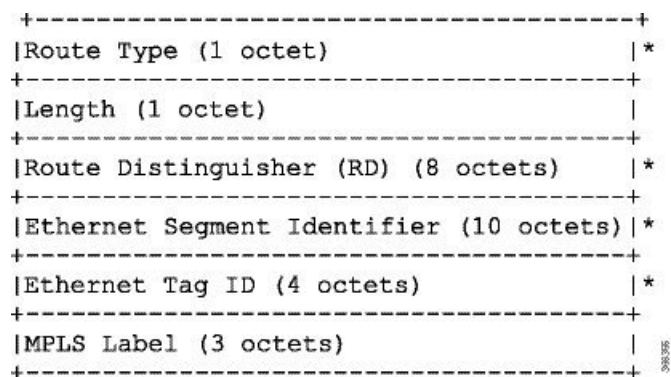
EVPN Route Types

The EVPN NLRI has the following different route types:

Route Type 1: Ethernet Auto-Discovery (AD) Route

The Ethernet (AD) routes are advertised on per EVI and per Ethernet Segment Identifier (ESI) basis. These routes are sent per Ethernet segment (ES). They carry the list of EVIs that belong to the ES. The ESI field is set to zero when a CE is single-homed.

An Ethernet A-D route type specific EVPN NLRI consists of the following fields:



NLRI Format: Route-type 1:

[Type] [Len] [RD] [ESI] [ETag] [MPLS Label]

Net attributes: [Type] [RD] [ESI] [ETag]

Path attributes: [MPLS Label]

Example

```
route-policy evpn-policy
  if rd in (1.1.1.1:0) [and/or evpn-route-type is 1] [and/or esi in (0a1.a2a3.a4a5.a6a7.a8a9)]
    [and/or etag is 4294967295] then
    set ..
  endif
end-policy
!
route-policy evpn-policy
  if rd in (1.1.1.2:0) [and/or evpn-route-type is 1] [and/or esi in
(00a1.a2a3.a4a5.a6a7.a8a9)] [and/or etag is 4294967295] then
    set ..
  endif
end-policy
```

Route Type 2: MAC/IP Advertisement Route

The host's IP and MAC addresses are advertised to the peers within NLRI. The control plane learning of MAC addresses reduces unknown unicast flooding.

A MAC/IP Advertisement Route type specific EVPN NLRI consists of the following fields:

Route Type (1 octet)	*
Length (1 octet)	
RD (8 octets)	*
Ethernet Segment Identifier (10 octets)	
Ethernet Tag ID (4 octets)	*
MAC Address Length (1 octet)	*
MAC Address (6 octets)	*
IP Address Length (1 octet)	*
IP Address (0, 4, or 16 octets)	*
MPLS Label1 (3 octets)	
MPLS Label2 (0 or 3 octets)	

3.063196

NLRI Format: Route-type 2:

[Type][Len][RD][ESI][ETag][MAC Addr Len][MAC Addr][IP Addr Len][IP Addr][MPLS Label1][MPLS Label2]

Net attributes: [Type][RD][ETag][MAC Addr Len][MAC Addr][IP Addr Len][IP Addr]

Path attributes: [ESI], [MPLS Label1], [MPLS Label2]

Example

```
route-policy evpn-policy
  if rd in (1.1.1.2:0) [and/or evpn-route-type is 2] [and/or esi in
(0000.0000.0000.0000)] [and/or etag is 0] [and/or macaddress in (0013.aabb.cddd)]
[and/or destination in (1.2.3.4/32)] then
    set ..
  endif
end-policy
```

Route Type 3: Inclusive Multicast Ethernet Tag Route

This route establishes the connection for broadcast, unknown unicast, and multicast (BUM) traffic from a source PE to a remote PE. This route is advertised on per VLAN and per ESI basis.

An Inclusive Multicast Ethernet Tag route type specific EVPN NLRI consists of the following fields:

Route Type (1 octet)	*
Length (1 octet)	
RD (8 octets)	*
Ethernet Tag ID (4 octets)	*
IP Address Length (1 octet)	*
Originating Router's IP Address (4 or 16 octets)	*

306357

NLRI Format: Route-type 3:

[Type][Len][RD][ETag][IP Addr Len][Originating Router's IP Addr]

Net attributes: [Type][RD][ETag][IP Addr Len][Originating Router's IP Addr]


Example

```
route-policy evpn-policy
  if rd in (1.1.1.1:300) [and/or evpn-route-type is 3] [and/or etag is 0] [and/or
evpn-originator in (1.1.1.1)] then
    set ..
  endif
end-policy
```

Route Type 4: Ethernet Segment Route

Ethernet segment routes enable to connect a CE device to two or PE devices. ES route enables the discovery of connected PE devices that are connected to the same Ethernet segment.

An Ethernet Segment route type specific EVPN NLRI consists of the following fields:



3-803139

NLRI Format: Route-type 4:

[Type][Len][RD][ESI][IP Addr Len][Originating Router's IP Addr]

Net attributes: [Type][RD][ESI][IP Addr Len][Originating Router's IP Addr]

Example

```
route-policy evpn-policy
  if rd in (1.1.1.1:0) [and/or evpn-route-type is 4] [and/or esi in
(00a1.a2a3.a4a5.a6a7.a8a9)] [and/or evpn-originator in (1.1.1.1)] then
    set ..
  endif
end-policy
```

Route Type 5: IP Prefix Route

An IP Prefix Route type specific EVPN NLRI consists of the following fields:

Route Type (1 octet)	*
Length (1 octet)	
RD (8 octets)	*
Ethernet Segment Identifier (10 octets)	
Ethernet Tag ID (4 octets)	*
IP Address Length (1 octet)	*
IP Address (4 or 16 octets)	*
GW IP Address (4 or 16 octets)	
MPLS Label (3 octets)	

NLRI Format: Route-type 5:

[Type][Len][RD][ESI][ETag][IP Addr Len][IP Addr][GW IP Addr][Label]

Net attributes: [Type][RD][ETag][IP Addr Len][IP Addr]

Path attributes: [ESI], [GW IP Addr], [Label]

Example

```
route-policy evpn-policy
  if rd in (30.30.30.30:1) [and/or evpn-route-type is 5] [and/or esi in
(0000.0000.0000.0000.0000)] [and/or etag is 0] [and/or destination in (12.2.0.0/16)] [and/or
evpn-gateway in (0.0.0.0)] then
    set ..
  endif
end-policy
```

EVPN RPL Attribute

Route Distinguisher

A Route Distinguisher (rd) attribute consists of eight octets. An rd can be specified for each of the EVPN route types. This attribute is not mandatory in route-policy.

Example

```
rd in (1.2.3.4:0)
```

EVPN Route Type

EVPN route type attribute consists of one octet. This specifies the EVPN route type. The EVPN route type attribute is used to identify a specific EVPN NLRI prefix format. It is a net attribute in all EVPN route types.

Example

```
evpn-route-type is 3
```

The following are the various EVPN route types that can be used:

```
1 - ethernet-ad
2 - mac-advertisement
3 - inclusive-multicast
4 - ethernet-segment
5 - ip-advertisement
```

IP Prefix

An IP prefix attribute holds IPv4 or IPv6 prefix match specification, each of which has four parts: an address, a mask length, a minimum matching length, and a maximum matching length. The address is required, but the other three parts are optional. When IP prefix is specified in EVPN route type 2, it represents either a IPv4 or IPv6 host IP Address (/32 or /128). When IP prefix is specified in EVPN route type 5, it represents either IPv4 or IPv6 subnet. It is a net attribute in EVPN route type 2 and 5.

Example

```
destination in (128.47.10.2/32)
destination in (128.47.0.0/16)
destination in (128:47::1/128)
destination in (128:47::0/112)
```

esi

An Ethernet Segment Identifier (ESI) attribute consists of 10 octets. It is a net attribute in EVPN route type 1 and 4, and a path attribute in EVPN route type 2 and 5.

Example

```
esi in (ffff.ffff.ffff.ffff.fff0)
```

etag

An Ethernet tag attribute consists of four octets. An Ethernet tag identifies a particular broadcast domain, for example, a VLAN. An EVPN instance consists of one or more broadcast domains. It is a net attribute in EVPN route type 1, 2, 3 and 5.

Example

```
etag in (10000)
```

mac

The mac attribute consists of six octets. This attribute is a net attribute in EVPN route type 2.

Example

```
mac in (0206.acb1.e806)
```

evpn-originator

The evpn-originator attribute specifies the originating router's IP address (4 or 16 octets). This is a net attribute in EVPN route type 3 and 4.

Example

```
evpn-originator in (1.2.3.4)
```

evpn-gateway

The evpn-gateway attribute specifies the gateway IP address. The gateway IP address is a 32-bit or 128-bit field (IPv4 or IPv6), and encodes an overlay next-hop for the IP prefixes. The gateway IP address field can be zero if it is not used as an overlay next-hop. This is a path attribute in EVPN route type 5.

Example

```
evpn-gateway in (1.2.3.4)
```

EVPN RPL Attribute Set

In this context, the term set is used in its mathematical sense to mean an unordered collection of unique elements. The policy language provides sets as a container for groups of values for matching purposes. Sets are used in conditional expressions. The elements of the set are separated by commas. Null (empty) sets are allowed.

prefix-set

A prefix-set holds IPv4 or IPv6 prefix match specifications, each of which has four parts: an address, a mask length, a minimum matching length, and a maximum matching length. The address is required, but the other three parts are optional. The prefix-set specifies one or more IP prefixes.

Example

```
prefix-set ip_prefix_set
14.2.0.0/16,
54.0.0.0/16,
12.12.12.0/24,
50:50::1:0/112
end-set
```


mac-set

The mac-set specifies one or more MAC addresses.

Example

```
mac-set mac_address_set
1234.2345.6789,
2345.3456.7890
end-set
```

esi-set

The esi-set specifies one or more ESI's.

Example

```
esi-set evpn_esi_set
1234.2345.3456.4567.5678,
1234.2345.3456.4567.5670
end-set
```

etag-set

The etag-set specifies one or more Ethernet tags.

Example

```
etag-set evpn_etag_set
10000,
20000
end-set
```

Configure EVPN RPL Feature

The following section describe how to configure mac-set, esi-set, evpn-gateway, and evpn-originator.

```
/* Configuring a mac-set and referring it in a route-policy (Attach point - neighbor-in) */
Router# configure
Router(config)# mac-set demo_mac_set
Router(config-mac)# 1234.ffff.aaa3,
Router(config-mac)# 2323.4444.ffff
Router(config-mac)# end-set
Router(config)# !
Router(config)# route-policy policy_use_pass_mac_set
Router(config-rpl)# if mac in demo_mac_set then
Router(config-rpl-if)# set med 200
Router(config-rpl-if)# else
Router(config-rpl-else)# set med 1000
Router(config-rpl-else)# endif
Router(config-rpl)# end-policy
Router(config)# commit
```

```

Router(config)# router bgp 100
Router(config-bgp)# address-family l2vpn evpn
Router(config-bgp-af)# !
Router(config-bgp-af)# neighbor 10.0.0.10
Router(config-bgp-nbr)# remote-as 8
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr-af)# route-policy policy_use_pass_mac_set in
Router(config-bgp-nbr-af)# commit

/* Configuring a esi-set and referring it in a route-policy (Attach point - neighbor-in) */
Router# configure
Router(config)# esi-set demo_esi
Router(config-esi)# ad34.1233.1222.ffff.44ff,
Router(config-esi)# ad34.1233.1222.ffff.6666
Router(config-esi)# end-set
Router(config)# !
Router(config)# route-policy use_esi
Router(config-rpl)# if esi in demo_esi then
Router(config-rpl-if)# set local-preference 100
Router(config-rpl-if)# else
Router(config-rpl-else)# set local-preference 300
Router(config-rpl-else)# endif
Router(config-rpl)# end-policy
Router(config)# commit

/* Configuring evpn-gateway/evpn-originator in a route-policy (Attach point - neighbor-in
and out) */
Router# configure
Router(config)# route-policy gateway_demo
Router(config-rpl)# if evpn-gateway in (10.0.0.0/32) then
Router(config-rpl-if)# pass
Router(config-rpl-if)# endif
Router(config-rpl)# end-policy
Router(config)# commit
Router(config)# route-policy originator_demo
Router(config-rpl)# if evpn-originator in (10.0.0.1/32) then
Router(config-rpl-if)# set local-preference 100
Router(config-rpl-if)# else
Router(config-rpl-else)# set med 200
Router(config-rpl-else)# endif
Router(config-rpl)# end-policy
Router(config)# commit
Router(config)# router bgp 100
Router(config-bgp)# address-family ipv4 unicast
Router(config-bgp-af)# !
Router(config-bgp-af)# neighbor 10.0.0.10
Router(config-bgp-nbr)# remote-as 8
Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy gateway_demo in
Router(config-bgp-nbr-af)# route-policy originator_demo out
Router(config-bgp-nbr-af)# commit

```

Running Configuration

```

/* Configuring a mac-set and referring it in a route-policy (Attach point - neighbor-in) */
mac-set demo_mac_set
  1234.ffff.aaa3,
  2323.4444.ffff
end-set
!
route-policy policy_use_pass_mac_set

```

```

        if mac in demo_mac_set then
            set med 200
        else
            set med 1000
        endif
    end-policy
!
router bgp 100
    address-family l2vpn evpn
    !
    neighbor 10.0.0.10
        remote-as 8
        address-family l2vpn evpn
        route-policy policy_use_pass_mac_set in
    !
!
end

/* Configuring a esi-set and referring it in a route-policy (Attach point - neighbor-in) */
Wed Oct 26 11:52:23.720 IST
esi-set demo_es1
    ad34.1233.1222.ffff.44ff,
    ad34.1233.1222.ffff.6666
end-set
!
route-policy use_es1
    if esi in demo_es1 then
        set local-preference 100
    else
        set local-preference 300
    endif
end-policy

```

EVPN Route Policy Examples

```

route-policy ex_2
    if rd in (2.2.18.2:1004) and evpn-route-type is 1 then
        drop
    elseif rd in (2.2.18.2:1009) and evpn-route-type is 1 then
        drop
    else
        pass
    endif
end-policy
!
route-policy ex_3
    if evpn-route-type is 5 then
        set extcommunity bandwidth (100:9999)
    else
        pass
    endif
end-policy
!
route-policy samp
end-policy
!
route-policy sampl
    if rd in (30.0.101.2:0) then
        pass
    endif
end-policy

```

```

!
route-policy samp2
  if rd in (30.0.101.2:0, 1:1) then
    pass
  endif
end-policy
!
route-policy samp3
  if rd in (*:*) then
    pass
  endif
end-policy
!
route-policy samp4
  if rd in (30.0.101.2:*) then
    pass
  endif
end-policy
!
route-policy samp5
  if evpn-route-type is 1 then
    pass
  endif
end-policy
!
route-policy samp6
  if evpn-route-type is 2 or evpn-route-type is 5 then
    pass
  endif
end-policy
!
route-policy samp7
  if evpn-route-type is 4 or evpn-route-type is 3 then
    pass
  endif
end-policy
!
route-policy samp8
  if evpn-route-type is 1 or evpn-route-type is 2 or evpn-route-type is 3 then
    pass
  endif
end-policy
!
route-policy samp9
  if evpn-route-type is 1 or evpn-route-type is 2 or evpn-route-type is 3 or evpn-route-type
  is 4 then
    pass
  endif
end-policy
!
route-policy test1
  if evpn-route-type is 2 then
    set next-hop 10.2.3.4
  else
    pass
  endif
end-policy
!
route-policy test2
  if evpn-route-type is 2 then
    set next-hop 10.10.10.10
  else
    drop
  endif
end-policy

```

```
end-policy
!
route-policy test3
  if evpn-route-type is 1 then
    set tag 9988
  else
    pass
  endif
end-policy
!
route-policy samp21
  if mac in (6000.6000.6000) then
    pass
  endif
end-policy
!
route-policy samp22
  if extcommunity rt matches-any (100:1001) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp23
  if evpn-route-type is 1 and esi in (aaaa.bbbb.cccc.dddd.eeee) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp24
  if evpn-route-type is 5 and extcommunity rt matches-any (100:1001) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp25
  if evpn-route-type is 2 and esi in (1234.1234.1234.1234.1236) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp26
  if etag in (20000) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp27
  if destination in (99.99.99.1) and etag in (20000) then
    pass
  else
    drop
  endif
end-policy
!
```

```
route-policy samp31
  if evpn-route-type is 1 or evpn-route-type is 2 or evpn-route-type is 3 or evpn-route-type
  is 4 or evpn-route-type is 5 then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp33
  if esi in evpn_esi_set1 then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp34
  if destination in (90:1:1::9/128) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp35
  if destination in evpn_prefix_set1 then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp36
  if evpn-route-type is 3 and evpn-originator in (80:1:1::3) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp37
  if evpn-gateway in (10:10::10) then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp38
  if mac in evpn_mac_set1 then
    pass
  else
    drop
  endif
end-policy
!
route-policy samp39
  if mac in (6000.6000.6002) then
    pass
  else
    drop
  endif
end-policy
```

```
!  
route-policy samp41  
  if evpn-gateway in (10.10.10.10, 10:10::10) then  
    pass  
  else  
    drop  
  endif  
end-policy  
!  
route-policy samp42  
  if evpn-originator in (24.162.160.1/32, 70:1:1::1/128) then  
    pass  
  else  
    drop  
  endif  
end-policy  
!  
route-policy example  
  if rd in (62300:1903) and evpn-route-type is 1 then  
    drop  
  elseif rd in (62300:19032) and evpn-route-type is 1 then  
    drop  
  else  
    pass  
  endif  
end-policy  
!  
route-policy samp100  
  if evpn-route-type is 4 or evpn-route-type is 5 then  
    drop  
  else  
    pass  
  endif  
end-policy  
!  
route-policy samp101  
  if evpn-route-type is 4 then  
    drop  
  else  
    pass  
  endif  
end-policy  
!  
route-policy samp102  
  if evpn-route-type is 4 then  
    drop  
  elseif evpn-route-type is 5 then  
    drop  
  else  
    pass  
  endif  
end-policy  
!  
route-policy samp103  
  if evpn-route-type is 2 and destination in evpn_prefix_set1 then  
    drop  
  else  
    pass  
  endif  
end-policy  
!  
route-policy samp104  
  if evpn-route-type is 1 and etag in evpn_etag_set1 then  
    drop
```

```

elseif evpn-route-type is 2 and mac in evpn_mac_set1 then
  drop
elseif evpn-route-type is 5 and esi in evpn_esi_set1 then
  drop
else
  pass
endif
end-policy
!
```

Support for DHCPv4 and DHCPv6 Client over BVI

The Support for DHCPv4 and DHCPv6 Client over the BVI feature allows you to configure DHCPv4 and DHCPv6 client on the Bridged Virtual Interface (BVI). You can configure a BVI, and request DHCP IPv4 or IPv6 address on the BVI. This allows your customer's device to have initial connectivity to your network without user intervention in the field. After the device is connected to your network, the customer devices can push a node-specific configuration with static IP addresses on a different BVI for customer deployment.

Configure DHCPv4 and DHCPv6 Client over BVI

Perform the following tasks to configure DHCPv4 and DHCPv6 client over BVI:

- Configure AC interface
- Configure L2VPN
- Configure BVI

Configuration Example

```

/* Configure AC interface */
Router# configure
Router(config)# interface tenGigE 0/5/0/1/1
Router(config-if)# bundle id 1 mode on
Router(config-if)# exit
Router(config)# interface Bundle-Ether1
Router(config-if)# no shut
Router(config-if)# exit
Router(config)# interface bundle-ether 1.100 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 100
Router(config-l2vpn-subif)# rewrite ingress tag pop 1 symmetric
Router(config-l2vpn-subif)# commit

/* Configure L2VPN */
Router # configure
Router(config)# l2vpn
Router(config-l2vpn)# bridge group BVI
Router(config-l2vpn-bg)# bridge-domain bvi
Router(config-l2vpn-bg-bd)# interface Bundle-Ether1.100
Router(config-l2vpn-bg-bd-ac)#exit
Router(config-l2vpn-bg-bd)# routed interface BVII1
Router(config-l2vpn-bg-bd-bvi)# commit

/* Configure BVI */
Router# configure
Router(config)# interface BVII1
```



```
Router(config-if)# ipv4 address dhcp
Router(config-if)# ipv6 address dhcp
Router(config-if)# commit
```

Running Configuration

This section shows the DHCPv4 and DHCPv6 client over BVI running configuration.

```
interface TenGigE0/5/0/1/1
bundle id 1 mode on
!
interface Bundle-Ether1
!
interface Bundle-Ether1.100 l2transport
encapsulation dot1q 100
rewrite ingress tag pop 1 symmetric
!
l2vpn
bridge group BVI
  bridge-domain bvi
    interface Bundle-Ether1.100
      !
      routed interface BVI1
      !
    !
  !
interface BVI1
ipv4 address dhcp
ipv6 address dhcp
!
```

Verification

The show output given in the following section display the details of DHCPv4 and DHCPv6 client over BVI configuration.

```
Router# show l2vpn bridge-domain
Legend: pp = Partially Programmed.
Bridge group: BVI, bridge-domain: bvi, id: 0, state: up, ShgId: 0, MSTi: 0
  Aging: 300 s, MAC limit: 64000, Action: none, Notification: syslog
  Filter MAC addresses: 0
  ACs: 2 (2 up), VFIs: 0, PWs: 0 (0 up), PBBs: 0 (0 up), VNIs: 0 (0 up)
  List of ACs:
    BVI, state: up, BVI MAC addresses: 1
    BE1.100, state: up, Static MAC addresses: 0
  List of Access PWs:
  List of VFIs:
  List of Access VFIs:
```

```
Router# show dhcp ipv4 client
```

Interface name	IP Address	Binding State	Lease Time Rem
BVI1	172.16.0.2	BOUND	3598 secs (00:59:58)

```
Router# show dhcp ipv6 client
```

Interface name	IPv6 Address	State	Lease Time Rem
BVI1	2000::1	BOUND	2591982

```
Router# show dhcp ipv4 client bvi1 detail
```

```
-----
Client Interface name       : BVI1
Client Interface handle    : 0x8804054
Client ChAddr              : 008a.9628.ac8a
Client ID                  : BVI1.00:8a:96:28:ac:8a
Client State               : BOUND
Client IPv4 Address (Dhcp) : 172.16.0.2
Client IPv4 Address Mask   : 255.240.0.0
Client Lease Time Allocated : 3600 secs (01:00:00)
Client Lease Time Remaining : 3571 secs (00:59:31)
Client Selected Server Address: 172.16.0.1
Client Next Hop Address    : 0.0.0.0
-----
```

```
Router# show dhcp ipv4 client BVI1 statistics
```

```
Client Interface name       : BVI1
-----
CLIENT COUNTER(s)         | VALUE
-----
Num discovers sent         :      44
Num requests sent         :       1
Num offers received       :       1
Num acks received         :       1
-----
```

```
Router# show dhcp ipv6 client
```

```
-----
Interface name             IPv6 Address           State                  Lease Time Rem
-----
BVI1                       2000::1                BOUND                  2591685
-----
```

```
Router# show dhcp ipv6 client statistics-all
```

```
Interface name             : BVI1
Interface handle           : 0x8804054
VRF                        : 0x60000000
-----
TYPE                       | TRANSMIT | RECEIVE | DROP |
-----
SOLICIT                    | 17 | 0 | 0 |
ADVERTISE                   | 0 | 1 | 0 |
REQUEST                     | 1 | 0 | 0 |
REPLY                       | 0 | 2 | 0 |
CONFIRM                     | 0 | 0 | 0 |
RENEW                       | 1 | 0 | 0 |
REBIND                      | 0 | 0 | 0 |
RELEASE                     | 0 | 0 | 0 |
RECONFIG                    | 0 | 0 | 0 |
INFORM                      | 0 | 0 | 0 |
-----
TIMER                       | STARTED | STOPPED | EXPIRED |
-----
INIT                        | 1 | 0 | 1 |
VBIND                      | 0 | 0 | 0 |
RENEW                       | 2 | 1 | 0 |
REBIND                      | 2 | 1 | 0 |
RETRANS                     | 19 | 3 | 16 |
-----
```

VALID | 2 | 1 | 0 |

Configure DHCPv6 Client Options

You can configure different DHCPv6 client options to differentiate between clients as required. Configure different DHCPv6 client options to differentiate how a DHCPv6 client communicates with a DHCPv6 server. The different DHCPv6 client options that you can configure are:

- **DUID:** If the DUID DHCPv6 client option is configured on an interface, DHCPv6 client communicates with the DHCPv6 server through the link layer address.
- **Rapid Commit:** If the Rapid Commit DHCPv6 client option is configured on an interface, DHCPv6 client can obtain configuration parameters from the DHCPv6 server through a rapid two-step exchange (solicit and reply) instead of the default four-step exchange (solicit, advertise, request, and reply).
- **DHCP Options:** The various other DHCPv6 options that can be configured on a DHCPv6 client are:
 - **Option 15:** Option 15 is also known as the User Class option and it is used by a DHCPv6 client to identify the type or category of users or applications it represents.
 - **Option 16:** Option 16 is also known as the Vendor ID option and it is used by a DHCPv6 a client to identify the vendor that manufactured the hardware on which the client is running.
 - **Option 23:** Option 23 is also known as the Domain name Server (DNS) option provides a list of one or more IPv6 addresses of DNS recursive name servers to which a client's DNS resolver can send DNS queries.
 - **Option 24:** Option 24 is also known as the Domain List option and it specifies the domain search list that the client uses to resolve hostnames with the DNS.
- **DHCP Timers:** This option is used to set different timer value for DHCP client configurations. The various DHCP timer options are:
 - **Release-timeout:** It is used to set retransmission timeout value for the initial release message.
 - **Req-max-rt:** It is used to set the maximum retransmission timeout value for the request message.
 - **Req-timeout:** It is used to set the initial request timeout value of the request message.
 - **Sol-max-delay:** It is used to set the maximum delay time of the first solicit message.
 - **Sol-max-rt:** It is used to set the maximum solicit retransmission time.
 - **Sol-time-out:** It is used to set the initial timeout value of the solicit message.

Configuration Example

Perform this task to configure DHCPv6 client options on a BVI interface.

```
Router# configure
Router(config)# interface BVI 10
Router(config-if)# ipv6 address dhcp-client-options
Router(config-dhcpv6-client)# duid linked-layer-address
Router(config-dhcpv6-client)# rapid-commit
Router(config-dhcpv6-client)# timers release-timeout 3
Router(config-dhcpv6-client)# timers sol-max-delay 1
Router(config-dhcpv6-client)# timers sol-time-out 1
```

```

Router(config-dhcpv6-client)# timers sol-max-rt 120
Router(config-dhcpv6-client)# timers req-max-rt 30
Router(config-dhcpv6-client)# timers req-timeout 1
Router(config-dhcpv6-client)# commit

```

Verification

To verify the DHCPv6 client options, use the **show dhcp ipv6 client BVI10 detail** command.

```

Router# show dhcp ipv6 client BVI10 detail
Wed Jun 10 16:19:21.272 IST

-----
Client Interface name : MgmtEth0/0/CPU0/1
Client Interface handle : 0x4040
Client MACAddr : 02f0.2b39.44be
Client State : BOUND
Client Link Local Address : fe80::f0:2bff:fe39:44be
Client IPv6 Address (Dhcp) : 600:1::12
Lease Remaining (in secs) : 74
DUID : 0003000102f02b3944be

Client Configuration
Timers
SOL_MAX_DELAY : 1 secs (00:00:01)
SOL_TIMEOUT : 1 secs (00:00:01)
SOL_MAX_RT : 120 secs (00:02:00)
REQ_TIMEOUT : 1 secs (00:00:01)
REQ_MAX_RT : 30 secs (00:00:30)
REL_TIMEOUT : 3 secs (00:00:01)

Options
RAPID-COMMIT : True
USER-CLASS : ciscoupnp
VENDOR-CLASS : vendor
DNS-SERVERS : True
DOMAIN-LIST : True

DUID Type : DUID_LL

Server Information
Server Address : fe80::d2:a1ff:feb2:3b9f
Preference : 0
DUID : 000300010206826e2e00
Status : SUCCESS
IA-NA
Status : SUCCESS
IAID : 0x40400001
T1 : 60 secs (00:01:00)
T2 : 96 secs (00:01:36)
IA-ADDR
IA NA Address : 600:1::12
Preferred Time : 120 secs (00:02:00)
Valid Time : 120 secs (00:02:00)
Flags : 0x0

```

Related Topics

- [Support for DHCPv4 and DHCPv6 Client over BVI, on page 126](#)

Associated Commands

- show l2vpn bridge-domain
- show dhcp ipv4 client
- show dhcp ipv6 client
- show dhcp ipv4 client bvi



CHAPTER 8

Configure EVPN IRB


This chapter introduces you to Ethernet VPN (EVPN) Integrated Routing and Bridging (IRB) feature and describe how you can configure the EVPN IRB feature.

- [EVPN IRB](#) , on page 133
- [EVPN Single-Homing Access Gateway](#) , on page 134
- [EVPN Multihoming All-Active](#), on page 135
- [Enable Auto-BGP RT with Manual ESI Configuration](#), on page 136
- [Supported EVPN IRB Scenarios](#), on page 136
- [Distributed Anycast Gateway](#), on page 136
- [VM Mobility Support](#), on page 139
- [Configuring EVPN IRB](#), on page 141
- [Running Configuration for EVPN IRB](#), on page 143
- [Verify EVPN IRB](#), on page 144
- [EVPN IPv6 Hosts with Mobility](#), on page 144

EVPN IRB

EVPN IRB feature enables a Layer 2 VPN and an Layer 3 VPN overlay that allows end hosts across the overlay to communicate with each other within the same subnet and across different subnets within the VPN.

Figure 20: EVPN IRB



The benefit of EVPN IRB is that it allows the hosts in an IP subnet to be provisioned anywhere in the data center. When a virtual machine (VM) in a subnet is provisioned behind a EVPN PE, and another VM is required in the same subnet, it can be provisioned behind another EVPN PE. The VMs do not have to be localized; they need not be directly connected; or be in the same complex. The VM is allowed to move across in the same subnet. Availability of IP MPLS network across all the EVPN PEs enables the provisioning of VM mobility. The EVPN PEs route traffic to each other through MPLS encapsulation.

The EVPN PEs are connected to each other by a spine so they have IP reachability to each other's loopback interfaces. The IP network and MPLS tunnels existing between these EVPN PEs constitute the IP MPLS underlay fabric.

You can configure the MPLS tunnels to tunnel Layer 2 traffic, and to overlay VPN on these tunnels. EVPN control plane distributes both Layer 2 MAC reachability and Layer 3 IP reachability for hosts within the context of the VPN; it overlays a tenant's VPN network on top of the MPLS underlay fabric. Thus you can have tenant's hosts, which are in the same subnet layer 2 domain, but distributed across the fabric, communicate to each other as if they are in a Layer 2 network.

The Layer 2 VLAN and the corresponding IP subnet are not only a network of physically connected hosts on Layer 2 links, but an overlaid network on top of underlayed IP MPLS fabric which is spread across the datacenter.

A routing service, which enables stretching of the subnet across the fabric, is available. It also provides Layer 3 VPN and performs routing between subnets within the context of the Layer 3 VPN. The EVPN PEs provide Layer 2 bridging service between hosts that are spread across the fabric within a Layer 2 domain that is stretched across the fabric, and Layer 3 VPN service or inter-subnet routing service for hosts in different subnets within Layer 3 VPN. For example, as shown in the above topology diagram, the two VM are in the same subnet but they are not connected directly through each other through a Layer 2 link. The Layer 2 link is replaced by MPLS tunnels that are connecting them. The whole fabric acts as a single switch and bridges traffic from one VM to the other. This also enables VM mobility.



Note Egress marking is not supported on L2 interfaces in a bridge domain.


In the above topology diagram, the VMs, VM1 and VM2 are connected each other. When VM2 migrates to a different switch and different server, the VM's current MAC address and IP address are retained. When the subnet is stretched between two EVPN PEs, the same IRB configuration is applied on both the devices.

For stretching within the same subnet, you must configure the AC interface and the EVI; it is not required to configure IRB interface or VRF.

EVPN Single-Homing Access Gateway

The EVPN provider edge (PE) devices learn the MAC address and IP address from the ARP traffic that they receive from the customer edge (CE) devices. The PEs create the MAC+IP routes. The PEs advertise the MAC+IP routes to MPLS core. They inject the host IP routes to IP-VPN gateway. Subnet routes are also advertised from the access EVPN PEs in addition to host routes. All the PE nodes add the host routes in the IP-VRF table. The EVPN PE nodes add MAC route to the MAC-VRF table. The IP-VPN PE advertise the subnet routes to the provider edge devices which add the subnet routes to IP-VRF table. On the PE devices, IRB gateway IP addresses and MAC addresses are not advertised through BGP. IRB gateway IP addresses or MAC addresses are used to send ARP requests towards the datacenter CEs.

Figure 21: EVPN Single-Homing Access Gateway




The above topology depicts how EVPN single-homing access gateway enables network connectivity by allowing a CE device to connect to one PE device. The PE device is attached to the Ethernet Segment through bundle or physical interfaces. Null Ethernet Segment Identifier (ESI) is used for single-homing.

EVPN Multihoming All-Active

In EVPN IRB, both EVPN and IP VPN (both VPNv4 and VPNv6) address families are enabled between routers and Data Center Interconnect (DCI) gateways. When Layer 2 (L2) stretch is not available in multiple data centers (DC), routing is established through VPNv4 or VPNv6 routes. When Layer 2 stretch is available, host routing is applied where IP-MAC routes are learnt by ARP and are distributed to EVPN/BGP. In remote peer gateway, these IP-MAC EVPN routes are imported into IP VPN routing table from EVPN route-type 2 routes with secondary label and Layer 3 VRF route-target.

Figure 22: EVPN Multi-Homing All-Active



The above topology describes how EVPN Multi-homing access gateway enables redundant network connectivity by allowing a CE device to connect to more than one PE device. Disruptions to the network connectivity are prevented by allowing a CE device to be connected to a PE device or several PE devices through multi-homing. Ethernet segment is the bunch of Ethernet links through which a CE device is connected to more than one PE devices. The All-Active Link Aggregation Group bundle operates as an Ethernet segment. Only MC bundles that operates between two chassis are supported.

Enable Auto-BGP RT with Manual ESI Configuration

Configuring an ES-Import RT was previously mandatory for Type 0 ESI. The ES-Import RT is auto-extracted by default, and the configuration serves to override the default value. This feature is based on [RFC 7432](#) but applied specifically to ESI Type 0. For more information, see Section 5 of [RFC 7432](#).

Supported EVPN IRB Scenarios

EVPN IRB supports the following scenarios:

- Dual-homing supports the following methods:
 - Only one EFP is supported per ESI per EVI
 - Only all-active mode is supported
 - Only two PE gateways in a redundancy group
- Single-homing supports the following methods:
 - Physical
 - VLAN
 - Bundle-ethernet
 - QinQ access
- Only IPv4 is supported.
- Subnet-stretch feature with EVPN IRB is only supported in VRF and is not supported in global VRF. In other words, EVPN IRB with EV-LAG multihoming is supported in global VRF without subnet being stretched beyond the multi-homing leafs

Distributed Anycast Gateway

EVPN IRB for the given subnet is configured on all the EVPN PEs that are hosted on this subnet. To facilitate optimal routing while supporting transparent virtual machine mobility, hosts are configured with a single default gateway address for their local subnet. That single (anycast) gateway address is configured with a single (anycast) MAC address on all EVPN PE nodes locally supporting that subnet. This process is repeated for each locally defined subnet requires Anycast Gateway support.

The host-to-host Layer 3 traffic, similar to Layer 3 VPN PE-PE forwarding, is routed on the source EVPN PE to the destination EVPN PE next-hop over an IP or MPLS tunnel, where it is routed again to the directly connected host. Such forwarding is also known as Symmetric IRB because the Layer 3 flows are routed at both the source and destination EVPN PEs.

The following are the solutions that are part of the Distributed Anycast Gateway feature:

EVPN IRB with All-Active Multi-Homing without Subnet Stretch or Host-Routing across the Fabric

For those subnets that are local to a set of multi-homing EVPN PEs, EVPN IRB Distributed Anycast Gateway is established through subnet routes that are advertised using EVPN Route Type 5 to VRF-hosting remote leafs. Though there is no need for the /32 routes within the subnet to be advertised, host MAC and ARP entries have to be synced across the EVPN PE to which the servers are multi-homed.

This type of multi-homing has the following characteristics:

- All-active EV LAG on access
- Layer 3 ECMP for the fabric for dual-homed hosts based on subnet routes
- Absence of Layer 2 subnet stretch over the fabric
- Layer 2 stretch within redundancy group of leafs with orphan ports

Prefix-routing solution for a non-stretched subnet is summarized as below:

Across multi-homing EVPN PEs:

- Local ARP cache and MAC addresses are synchronized for dual-homed hosts through EVPN MAC+IP host route advertisements. They are imported as local, and are based on the local ESI match, for optimal forwarding to the access gateway.
- Orphan MAC addresses and host IP addresses are installed as remote addresses over the fabric.
- ES/EAD routes are exchanged for the designated forwarder (DF) election and split-horizon label.

Across remote EVPN PEs:

- Dual-homed MAC+IP EVPN Route Type 2 is exchanged with the ESI, EVI Label, Layer 2-Route Type. It is not imported across the fabric, if there is no subnet stretch or host-routing.
- The subnet IP EVPN Route Type 5 is exchanged with VRF label and Layer 3-Route Type.
- Layer 3 Route Type for the VRFs is imported that are present locally.
- Layer 2 Route Type for locally present BDs is imported. It is only imported from the leaf in the same redundancy group, if BD is not stretched.

EVPN IRB with All-Active Multihoming with Subnet Stretch or Host-Routing across the Fabric

For a bridge domain or subnet that is stretched across remote EVPN PEs, both /32 host routes and MAC routes are distributed in a EVPN overlay control plane to enable Layer 2 and Layer 3 traffic to the end points in a stretched subnet.

This type of multihoming has the following characteristics:

- All-active EV-LAG on the access gateway
- Layer 2 or Layer 3 ECMP for the fabric for dual-homed hosts based on Route Type 1 and Route Type 2

- Layer 3 unipath over the fabric for single-homed hosts based on Route Type 2
- Layer 2 subnet stretch over the fabric
- Layer 2 stretch within redundancy group of leafs with orphan ports

MAC and host routing solution for a stretched subnet is summarized as follows:

Across multihoming EVPN PEs:

- The Local ARP cache and MAC addresses are synchronized for dual-homed hosts through EVPN MAC+IP host route advertisements. They are imported as local, based on the local ESI match, for optimal forwarding to the access gateway.
- Synchronized MAC+IP are re-originated for inter-subnet Layer 3 ECMP.
- Orphan MAC address and host IP address are installed as remote addresses over the fabric.
- ES/EAD route is exchanged for designated forwarder (DF) election and split-horizon label.

Across remote EVPN PEs:

- Dual-homed MAC+IP EVPN Route Type 2 is exchanged with ESI, EVI label, Layer 2-Route Type, VRF label, and Layer 3-Route Type.
- Subnet IP EVPN Route Type 5 is exchanged for VRF label, Layer 3-Route Type for silent hosts, and non-stretched subnets.
- Layer 3 Route Type is imported for locally present VRFs.
- Layer 2 Route Type is imported for locally present bridge domains.

MAC and IP Unicast Control Plane

This use case has following types:

Prefix Routing or No Subnet Stretch

IP reachability across the fabric is established using subnet prefix routes that are advertised using EVPN Route Type 5 with the VPN label and VRF RTs. Host ARP and MAC sync are established across multi-homing EVPN PEs using MAC+IP Route Type 2 based on a shared ESI to enable local switching through both the multi-homing EVPN PEs.

Host Routing or Stretched Subnet

When a host is discovered through ARP, the MAC and IP Route Type 2 is advertised with both MAC VRF and IP VRF router targets, and with VPN labels for both MAC-VRF and IP-VRF. Particularly, the VRF route targets and Layer 3 VPN label are associated with Route Type 2 to achieve PE-PE IP routing identical to traditional L3VPNs. A remote EVPN PE installs IP/32 entries directly in Layer 3 VRF table through the advertising EVPN PE next-hop with the Layer 3 VPN label encapsulation, much like a Layer 3 VPN imposition PE. This approach avoids the need to install separate adjacency rewrites for each remote host in a stretched subnet. Instead, it inherits a key Layer 3 VPN scale benefit of being able to share a common forwarding rewrite or load-balance resource across all IP host entries reachable through a set of EVPN PEs.

ARP and MAC sync

For hosts that are connected through LAG to more than one EVPN PE, the local host ARP and MAC entries are learnt in data plane on either or both of the multihoming EVPN PEs. Local ARP and MAC entries are

synced across the two multihoming EVPN PEs using MAC and IP Route Type 2 based on a shared ESI to enable local switching through both the multihoming EVPN PEs. Essentially, a MAC and IP Route Type 2 that is received with a local ESI causes the installation of a synced MAC entry that points to the local AC port, and a synced ARP entry that is installed on the local BVI interface.



Note Only one Ethernet Flow Point (EFP) is supported per non-Zero ESI per bridge domain or EVI. This is a limitation of EVPN.

MAC and IP Route Re-origination

MAC and IP Route Type 2 received with a local ESI, which is used to sync MAC and ARP entries, is also re-originated from the router that installs a SYNC entry, if the host is not locally learnt and advertised based on local learning. This route re-origination is required to establish overlay IP ECMP paths on remote EVPN PEs, and to minimize traffic hit on local AC link failures, that can result in MAC and IP route withdraw in the overlay.



Note If custom or static MAC address is configured on a BVI interface, the MAC address on the wire may be different than what is configured. This has no operational or functional impact.

Intra-subnet Unicast Data Plane

The Layer 2 traffic is bridged on the source EVPN PE using ECMP paths to remote EVPN PEs, established through MAC+IP RT2, for every ES and for every EVI, ES and EAD Route Type 2 routes that are advertised from the local EVPN PEs.

Inter-subnet Unicast Data Plane

Inter-subnet traffic is routed on the source ToRs through overlay ECMP to the destination ToR next-hops. Data packets are encapsulated with the VPN label advertised from the ToR and tunnel label for the BGP next-hop towards the spine. It is then routed again on the destination ToR using a local ARP adjacency towards the host. IP ECMP on the remote ToRs is established through local and re-originated routes advertised from the local ToRs.

VM Mobility Support

VM mobility is the ability of virtual machines to migrate between one server and another while retaining their existing MAC and IP addresses.

The following are the two key components in EVPN Route Type 2 that enable VM Mobility:

- Host MAC advertisement component that is imported into local bridge MAC table, and Layer 2 bridged traffic across the network overlay.
- Host IP advertisement component that is imported into the IP routing table in a symmetric IRB design, enables routed traffic across the network overlay.

The above-mentioned components are advertised together in a single MAC + IP host route advertisement. An additional MAC-only route could also be advertised.

The following behaviors of VM are supported. The VM can:

- retain existing MAC and acquire a new IP address
- retain existing IP address and acquire a new MAC
- retain both existing MAC and IP address

MAC and MAC-IP Sequence Numbers

The IRB gateway device assigns, manages, and advertises sequence numbers that are associated with the locally learnt MAC routes through hardware learning, and the locally learnt MAC-IP routes through ARP.

Synchronized MAC and MAC-IP Sequence Numbers

In a host that is multi-homed to two ToRs, the locally learnt MAC and MAC-IP routes are synchronized across the two multi-homing peers through Route Type 2 learnt routes with a local ESI. So a device could have either MAC and MAC-IP, or both of them, learnt through both synchronized and local learning. Sequence numbers are synchronized across local and synchronized routes, because of which the sequence number that is advertised from the two ToRs for a given route is always the same. In certain situations, remote-sync route with same ESI can have a higher sequence number than a local route. In such a case, the local route sequence number is bumped up to match remote-sync route sequence number.

Local Sequence Number Updates

Host mobility is triggered when a local route is learnt while a remote route already exists. When mobility occurs, the local route is assigned a sequence number that is one higher than the existing remote route. This new local route is then advertised to the rest of the network.

Best Route Selection after Host Movement

When a host moves, the EVPN-PE at the new location of the host generates and advertises a higher sequence route to the network. When a higher sequence number route is received, as per RFC 7432, it is considered as the new best route and it is used for forwarding traffic. Best route selection is done for both MAC and MAC-IP routes.

Stale Route Deletion after a Host Movement

After a host moves from local to remote ESI, if a remote route from a different ESI is received and if a local route for the same host with a lower sequence number exists, then the local route is deleted and is withdrawn from the network.

The new higher sequence number remote MAC route is now considered best and is used to forward traffic. An ARP probe is sent to the host at the old local location. Because the host is at new remote location, probe will not succeed, resulting in clearing old local MAC-IP route.

Host Movement Detection through GARP

If a host sends a Gratuitous ARP (GARP) at its new location after a movement, the local MAC and local MAC-IP learning independently trigger mobility for both routes.

Host Move Detection with Silent Host

If a host does not send a GARP or a data packet at its new location following a move, the aging of the local MAC at the old location triggers mobility for both routes.

Host Move Detection without GARP with Data Packet

If the host does not send a GARP following a move, a data packet from the host triggers a proactive ARP probe to discover host MAC-IP and trigger mobility for this host across the overlay.

Duplicate MAC Detection

Duplicate MAC detection and freezing is supported as per RFC 7432.

Detection: Duplicate detection and recovery parameters are configurable. The default configuration is five times in 180 seconds and route freezing after three duplicate cycles. With the default configuration, when a host moves five times in 180 seconds, it is marked as duplicate for 30 seconds. Route advertisement for hosts in Duplicate state is suppressed. Host is taken out of duplicate state after 30 seconds. After a host is detected as duplicate for 3 times, on the fourth duplicate cycle, the host is permanently frozen. All route advertisements are suppressed for the frozen hosts.

In multi-homed hosts, a MAC is not necessarily learnt locally but is learnt through synchronization. Duplicate detection is supported for both local and remote-sync hosts. Remote-sync routes are differentiated from remote routes.

MAC-IP Handling: If the MAC route is in duplicate or frozen state, the corresponding local MAC-IP is updated, except that the route deletes are not withheld.

Duplicate State Handling: When a host is in duplicate state, route advertisements are suppressed. However, local routes are programmed in hardware so that traffic on local EVPN-PE is forwarded to the local host.

Recovery: It is possible to unfreeze permanently frozen hosts. The following is the recommended procedure to clear frozen hosts:

- Shutdown the host which is causing duplicate traffic.
- Use the **clear l2route evpn frozen-mac frozen-flag** command to clear the frozen hosts.

Configuring EVPN IRB

```
/* Configure CEF to prefer RIB prefixes over adjacency prefixes.*/
```

```
RP/0/RSP0/CPU0:router# configure  
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether 3  
RP/0/RSP0/CPU0:router(config-if)# l2cp system mac 1.1.1
```

```

RP/0/RSP0/CPU0:router(config-if)# exit
RP/0/RSP0/CPU0:router(config)# cef adjacency route override rib

/* Configure EVPN L3VRF per DC tenant. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# vrf irbl
RP/0/RSP0/CPU0:router(config-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-vrf-af)# import route-target 1000:1
RP/0/RSP0/CPU0:router(config-vrf-af)# export route-target 1000:1
RP/0/RSP0/CPU0:router(config-vrf-af)# exit

/* Configure Layer 2 attachment circuit (AC) from multichassis (MC) bundle interface, and
bridge-group virtual interface (BVI) per bridge domain. */
/* Note: When a VM migrates from one subnet to another (subnet stretching), apply the
following IRB configuration to both the EVPN PEs. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface bvi 1001
RP/0/RSP0/CPU0:router(config-if)# host-routing
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 10.10.0.4 255.255.255.0
RP/0/RSP0/CPU0:router(config-if)# ipv4 address 172.16.0.1 secondary
RP/0/RSP0/CPU0:router(config-if)# mac-address 2001:DB8::1

/* Configure EVPN Layer 2 bridging service. Note: This configuration is performed in Layer
2 gateway or bridging scenario. */

Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# bridge group 1
Router(config-l2vpn-bg)# bridge-domain 1-1
Router(config-l2vpn-bg-bd)# interface GigabitEthernet 0/0/0/1.1
Router(config-l2vpn-bg-bd-ac)# evi 1
Router(config-l2vpn-bg-bd-ac-evi)# commit
Router(config-l2vpnbg-bd-ac-evi)# exit

/* Configure BGP. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# router bgp 3107 router-id 192.168.1.1
RP/0/RSP0/CPU0:router(config-bgp)# vrf irbl
RP/0/RSP0/CPU0:router(config-bgp-vrf)# rd auto
RP/0/RSP0/CPU0:router(config-bgp-vrf)# address-family ipv4 unicast
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)# redistribute connected
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)# redistribute static
RP/0/RSP0/CPU0:router(config-bgp-vrf-af)# exit

/* Configure EVPN, and configure main bundle ethernet segment parameters in EVPN. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# evpn
RP/0/RSP0/CPU0:router(config-evpn)# evi 2001
RP/0/RSP0/CPU0:router(config-evpn-instance)# bgp
RP/0/RSP0/CPU0:router(config-evpn-evi-bgp)# route-target import 1000:1
RP/0/RSP0/CPU0:router(config-evpn-evi-bgp)# route-target export 1000:1

RP/0/RSP0/CPU0:router(config-evpn-evi)# advertise-mac
RP/0/RSP0/CPU0:router(config-evpn-evi)# unknown-unicast-suppression

```



```

/* Configure Layer 2 VPN. */

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# l2vpn
RP/0/RSP0/CPU0:router(config-l2vpn)# bridge group irb
RP/0/RSP0/CPU0:router(config-l2vpn-bg)# bridge-domain irb1
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd)# interface bundle-Ether3.1001
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-ac)# routed interface BVI100
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-bvi)# split-horizon group core
RP/0/RSP0/CPU0:router(config-l2vpn-bg-bd-bvi)# evi 10001

```

Running Configuration for EVPN IRB

```

/* Configure LACP */

interface Bundle-Ether3
  lacp system mac 1.1.1
!

/* Configure CEF adjacency overwrite. */

cef adjacency route override rib

/* Configure EVPN Layer 3 VRF per DC tenant. */

vrf irb1
address-family ipv4 unicast
  import route-target
    1000:1
  !
  export route-target
    1000:1
  !
!
!

/* Configure Layer 2 attachment circuit (AC) from multichassis (MC) bundle interface, and
bridge-group virtual interface (BVI) per bridge domain.*/

interface Bundle-Ether3.1001 l2transport
  encapsulation dot1q 1001
  rewrite ingress tag pop 1 symmetric
!
interface BVI1001
  host-routing
  vrf irb1
  ipv4 address 10.0.1.1 255.255.255.0
  mac-address 0000.3030.1
!

/* Configure BGP. */

router bgp 3107
  vrf irb1
  rd auto
  address-family ipv4 unicast

```

```

    redistribute connected
    redistribute static
    !
    !

/* Configure EVPN. */

evpn
evi 10001
    bgp
        route-target import 1000:1
        route-target export 1000:1
    !
    advertise-mac
    unknown-unicast-suppression
    !

/* Configure Layer2 VPN. */

l2vpn
bridge group irb
    bridge-domain irb1
        interface Bundle-Ether3.1001
        !
        routed interface BVI1001
            split-horizon group core
        !
        evi 10001
        !
    !

```

Verify EVPN IRB

EVPN IPv6 Hosts with Mobility

EVPN IPv6 Hosts with Mobility feature enables you to provide EVPN IPv6 service over IPv4-MPLS core network. This feature supports all-active multihoming and virtual machine (VM) or host move.

Service Providers (SPs) use a stable and established core with IPv4-MPLS backbone for providing IPv4 VPN services. The IPv6 VPN Provider Edge Transport over MPLS (IPv6 on Provider Edge Routers [6PE] and IPv6 on VPN Provider Edge Routers [6VPE]) facilitates SPs to offer IPv6 VPN services over IPv4 backbone without an IPv6 core. The provide edge (PE) routers run MP-iBGP to advertise IPv6 reachability and IPv6 label distribution. For 6PE, the labels are allocated per IPv6 prefix learnt from connected customer edge (CE) routers and for 6VPE, the PE router can be configured to allocate labels on a per-prefix or per-CE and per-VRF level.

Mobility Support

In global VRF, mobility is not supported. However, you can move a host from one ES to another ES within the same bridge domain. The host gets a new MAC address and IP address. The host can have multiple IP addresses for the same MAC address.

In non-default VRF, mobility is supported with the following conditions:

- Basic MAC move: The IP address and MAC address remains the same. You can move a host from one ES to another ES with the same IP address and MAC address

- Same MAC address but with a different IP address: The host gets a new IP address
- Same IP address but with a different MAC address: The host gets a new MAC address but retains the same IP address
- Multiple IP addresses with the same MAC address: Many VMs are involved in the same the MAC move

Restrictions

- In customer VRFs, when host routing is not configured, MAC-IP advertisement is different between zero ESI and non-zero ESI. When host routing is not configured, MAC-IP with non-zero ESI is advertised without L3 RT (VRF RT). MAC-IP with zero ESI is not advertised. The following table lists the behavior of MAC-IP advertisement with respect to ESI and host routing.

ESI Type	With host routing	Without host routing
MAC-IP with non-zero ESI	Advertised with L3 VRF RT	Advertised without L3 VRF RT
MAC-IP with zero ESI	Advertised with L3 VRF RT	Not advertised

- In global VRF, Layer 2 stretch is not supported.
- MAC move in global VRF is only supported if the host is within the same bridge domain. You can move a host from one ES to another ES within the same bridge domain.
- Duplication of IP address detection is not supported.
- Maximum number of leafs allowed per ESI is two.

Configure EVPN IPv6 Hosts with Mobility

Perform the following tasks to configure EVPN IPv6 Hosts with Mobility feature:

- Configure VRF
- Configure ISIS
- Configure BGP
- Configure AC interface
- Configure BVI interface
- Configure EVPN
- Configure L2VPN

**Note**

- You cannot configure the EVPN remote peer using the VPNv4 unicast if you have configured the **advertise vpnv4 unicast re-originated** command under the L2VPN EVPN address-family. You can either configure the VPNv4 unicast or the advertise vpnv4 unicast re-originated under L2VPN EVPN address-family.
- You cannot configure the EVPN remote peer using the VPNv6 unicast if you have configured the **advertise vpnv6 unicast re-originated** command under the L2VPN EVPN address-family. You can either configure the VPNv6 unicast or the advertise vpnv6 unicast re-originated under L2VPN EVPN address-family.

```

/* Configure VRF */

Router# configure
Router(config)# vrf cust102
Router(config-vrf)# address-family ipv4 unicast
Router(config-vrf-af)# import route-target 160102:16102
Router(config-vrf-af)# export route-target 160102:16102
Router(config-vrf-af)# exit
!
Router(config-vrf)# address-family ipv6 unicast
Router(config-vrf-af)# import route-target 6160102:16102
Router(config-vrf-af)# export route-target 6160102:16102
Router(config-vrf-af)# commit
!

/* Configure ISIS */

Router# configure
Route(config)# router isis v6
Route(config-isis)# 49.0001.0000.0160.0005.00
Route(config-isis)# nsr
Route(config-isis)# log adjacency changes
Route(config-isis)# lsp-gen-interval maximum-wait 5000 initial-wait 1 secondary-wait
20
Route(config-isis)# lsp-mtu 1468
Route(config-isis)# lsp-refresh-interval 65000
Route(config-isis)# max-lsp-lifetime 65535
Route(config-isis)# address-family ipv4 unicast
Route(config-isis-af)# metric-style wide
Route(config-isis-af)# microloop avoidance protected
Route(config-isis-af)# spf-interval maximum-wait 5000 initial-wait 1 secondary-wait 20
Route(config-isis-af)# segment-routing mpls sr-prefer
Route(config-isis-af)# segment-routing prefix-sid-map advertise-local
Route(config-isis-af)# exit
!
Route(config-isis)# interface Bundle-Ether10
Route(config-isis-if)# point-to-point
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-af)# fast-reroute per-prefix
Route(config-isis-af)# fast-reroute per-prefix ti-lfa
Route(config-isis-af)# metric 10
Route(config-isis-af)# exit
!
Route(config-isis)# interface Bundle-Ether20

```

```

Route(config-isis-if) # point-to-point
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-af) # fast-reroute per-prefix
Route(config-isis-af) # fast-reroute per-prefix ti-lfa
Route(config-isis-af) # metric 10
Route(config-isis-af) # exit
!
Route(config-isis) # interface loopback0
Route(config-isis-if) # passive
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-af) # exit
!
Route(config-isis) # interface loopback10
Route(config-isis-if) # passive
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-af) # prefix-sid index 1605
Route(config-isis-af) # commit
Route(config-isis-af) # exit
!

/* Configure Segment Routing */

Router# configure
Router(config) # segment-routing
Router(config-sr) # global-block 16000 23999
Router(config-sr) # commit

/* Configure BGP */

Router(config) # router bgp 100
Router(config-bgp) # bfd minimum-interval 50
Router(config-bgp) # bfd multiplier 3
Router(config-bgp) # bgp router-id 160.0.0.5
Router(config-bgp) # address-family ipv4 unicast ---> To support V4 Global VRF
Router(config-bgp-af) # maximum-paths ibgp 10 unequal-cost ---> ECMP
Router(config-bgp-af) # redistribute connected --> V4 Global VRF
Router(config-bgp-af) # exit
!
Router(config-bgp) # address-family ipv4 unicast ---> VRF
Router(config-bgp-af) # vrf all
Router(config-bgp-af) # label mode per-vrf
Router(config-bgp-af) # exit
!
Router(config-bgp) # address-family ipv6 unicast ---> For 6PE
Router(config-bgp-af) # label mode per-vrf
Router(config-bgp-af) # maximum-paths ibgp 8
Router(config-bgp-af) # redistribute static
Router(config-bgp-af) # allocate-label all
Router(config-bgp-af) # exit
!
Router(config-bgp) # address-family vpnv6 unicast ---> 6 VPE
Router(config-bgp-af) # vrf all
Router(config-bgp-af) # label mode per-vrf
Router(config-bgp-af) # exit
!
Router(config-bgp) # address-family l2vpn evpn ----> EVPN
Router(config-bgp-af) # bgp implicit-import ----> Global VRF
Router(config-bgp-af) # exit
!
Router(config-bgp) # neighbor-group evpn-rr
Router(config-bgp-nbr) # remote-as 100
Router(config-bgp-nbr) # bfd fast-detect
Router(config-bgp-nbr) # update-source loopback0

```

```

Router(config-bgp-nbr)# address-family ipv4 unicast
Router(config-bgp-nbr-af)# route-policy pass-all in
Router(config-bgp-nbr-af)# route-policy nh-lo10 out
Router(config-bgp-nbr-af)# exit
!
Router(config-bgp-nbr)# address-family ipv6 labeled-unicast ----> For 6PE
Router(config-bgp-nbr-af)# route-policy pass-all out
Router(config-bgp-nbr-af)# exit
!
Router(config-bgp-nbr)# address-family l2vpn evpn
Router(config-bgp-nbr-af)# route-policy pass-all in
Router(config-bgp-nbr-af)# route-policy nh-lo10 out
Router(config-bgp-nbr-af)# advertise vpnv4 unicast re-originated -> For Route Type 5
Router(config-bgp-nbr-af)# advertise vpnv6 unicast re-originated -> For Route Type 5
Router(config-bgp-nbr-af)# exit
!
Router(config-bgp)# neighbor 160.0.0.1
Router(config-bgp-nbr)# use neighbor-group evpn-rr
Router(config-bgp-nbr)# exit
!
Router(config-bgp)# neighbor 160.0.0.2
Router(config-bgp-nbr)# use neighbor-group evpn-rr
Router(config-bgp-nbr)# exit
!
Router(config-bgp)# vrf all
Router(config-bgp-vrf)# rd 1605:102
Router(config-bgp-vrf-af)# address-family ipv4 unicast
Router(config-bgp-vrf-af)# label mode per-vrf
Router(config-bgp-vrf-af)# maximum-paths ibgp 10 unequal-cost
Router(config-bgp-vrf-af)# redistribute connected ----> Triggers Route Type 5
Router(config-bgp-vrf-af)# exit
!
Router(config-bgp-vrf)# address-family ipv6 unicast
Router(config-bgp-vrf-af)# label mode per-vrf
Router(config-bgp-vrf-af)# maximum-paths ibgp 10 unequal-cost
Router(config-bgp-vrf-af)# redistribute connected
Router(config-bgp-vrf-af)# exit
!

/* Configure AC interface */

Router(config)# interface Bundle-Ether1.102 l2transport
Router(config-l2vpn-subif)# encapsulation dot1q 102
Router(config-l2vpn-subif)# rewrite ingress tag pop 1 symmetric
Router(config-l2vpn-subif)# commit
Router(config-l2vpn-subif)# exit

/* Configure BVI interface */

Router(config)# interface BVI100
Router(config-if)# ipv4 address 56.78.100.1 255.255.255.0
Router(config-if)# ipv6 address 56:78:100::1/64
Router(config-if)# mac-address 22.22.22
Router(config-if)# exit
!
Router(config)# interface BVI102
Router(config-if)# host-routing
Router(config-if)# vrf cust102
Router(config-if-vrf)# ipv4 address 56.78.102.1 255.255.255.0
Router(config-if-vrf)# ipv6 address 56:78:100::1/64
Router(config-if-vrf)# ipv6 address 56:78:102::1/64
Router(config-if-vrf)# mac-address 22.22.22
Router(config-if)# commit

```

```

/* Configure CEF */ [Required for dual homing]

Router# configure
Router(config)# cef adjacency route override rib

/* Configure EVPN, and configure main bundle ethernet segment parameters in EVPN */

Router# configure
Router(config)# evpn
Router(config-evpn)# evi 102
Router(config-evpn-evi)# bgp
Router(config-evpn-evi)# rd 1605:102
Router(config-evpn-evi-bgp)# route-target import 160102:102
Router(config-evpn-evi-bgp)# route-target export 160102:102
Router(config-evpn-evi-bgp)# exit
Router(config-evpn-evi)# advertise-mac
Router(config-evpn-evi)# exit
!
Router(config-evpn)# interface Bundle-Ether1
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 56.56.56.56.56.56.56.01
Router(config-evpn-ac-es)# exit
!
Router(config-evpn)# interface Bundle-Ether2
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 56.56.56.56.56.56.56.02
Router(config-evpn-ac-es)# commit

/* Configure L2VPN */

Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# bridge group bg102
Router(config-l2vpn-bg)# bridge-domain bd102
Router(config-l2vpn-bg-bd)# interface Bundle-Ether1.102
Router(config-l2vpn-bg-bd-ac)# exit
!
Router(config-l2vpn-bg-bd)# interface Bundle-Ether2.102
Router(config-l2vpn-bg-bd-ac)# exit
!
Router(config-l2vpn-bg-bd)# interface Bundle-Ether3.102
Router(config-l2vpn-bg-bd-ac)# exit
!
Router(config-l2vpn-bg-bd)# interface Bundle-Ether4.102
Router(config-l2vpn-bg-bd-ac)# exit
!
Router(config-l2vpn-bg-bd)# interface Bundle-Ether5.102
Router(config-l2vpn-bg-bd-ac)# routed interface BVI102
Router(config-l2vpn-bg-bd-bvi)# evi 102
Router(config-l2vpn-bg-bd-bvi-evi)# commit

```

Running Configuration

```

/* Configure VRF */

vrf cust102
address-family ipv4 unicast
import route-target
160102:16102
!

```

```

export route-target
160102:16102
!
!
address-family ipv6 unicast
import route-target
6160102:16102
!
export route-target
6160102:16102
!
!
!

/ * Configure ISIS */

router isis v6
net 49.0001.0000.0160.0005.00
nsr
log adjacency changes
lsp-gen-interval maximum-wait 5000 initial-wait 1 secondary-wait 20
lsp-mtu 1468
lsp-refresh-interval 65000
max-lsp-lifetime 65535
address-family ipv4 unicast
metric-style wide
microloop avoidance protected
spf-interval maximum-wait 5000 initial-wait 1 secondary-wait 20
segment-routing mpls sr-prefer
segment-routing prefix-sid-map advertise-local
!
interface Bundle-Ether10
point-to-point
address-family ipv4 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
metric 10
!
!
interface Bundle-Ether20
point-to-point
address-family ipv4 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
metric 10
!
!
interface Loopback0
passive
address-family ipv4 unicast
!
!
interface Loopback10
passive
address-family ipv4 unicast
prefix-sid index 1605
!
!
!

/ * Configure Segment Routing */

segment-routing
global-block 16000 23999

```



```

!

/ * Configure BGP */

router bgp 100
  bfd minimum-interval 50
  bfd multiplier 3
  bgp router-id 160.0.0.5
  address-family ipv4 unicast      ---> To support V4 Global VRF
    maximum-paths ibgp 10 unequal-cost ---> ECMP
    redistribute connected        --> V4 Global VRF
  !
  address-family vpnv4 unicast ---> VRF
    vrf all
    label mode per-vrf
  !
  address-family ipv6 unicast      ---> For 6PE
    label mode per-vrf
    maximum-paths ibgp 8
    redistribute connected
    redistribute static
    allocate-label all
  !
  address-family vpnv6 unicast      ---> 6VPE
    vrf all
    label mode per-vrf
  !
  address-family l2vpn evpn        ----> EVPN
  bgp implicit-import              ----> Global VRF
  !

neighbor-group evpn-rr
  remote-as 100
  bfd fast-detect
  update-source Loopback0
  address-family ipv4 unicast
    route-policy pass-all in
    route-policy nh-lo10 out
  !
  address-family ipv6 labeled-unicast ----> For 6PE
  route-policy pass-all out
  !
  address-family l2vpn evpn
  route-policy pass-all in
  route-policy nh-lo10 out
  advertise vpnv4 unicast re-originated ---> For Route Type 5
  advertise vpnv6 unicast re-originated ----> For Route Type 5
  !
  !
  neighbor 160.0.0.1
  use neighbor-group evpn-rr
  !
  neighbor 160.0.0.2
  use neighbor-group evpn-rr
  !
  vrf cust102
  rd 1605:102
  address-family ipv4 unicast
  label mode per-vrf
  maximum-paths ibgp 10 unequal-cost
  redistribute connected <----- Triggers Route Type 5
  !
  address-family ipv6 unicast
  label mode per-vrf

```

```

maximum-paths ibgp 10 unequal-cost
redistribute connected
!
!

/* Configure AC interface */

interface Bundle-Ether1.102 l2transport
encapsulation dot1q 102
rewrite ingress tag pop 1 symmetric
!
/* Configure BVI interface */
interface BVI100
ipv4 address 56.78.100.1 255.255.255.0
ipv6 address 56:78:100::1/64
mac-address 22.22.22
!
interface BVI102
host-routing
vrf cust102
ipv4 address 56.78.102.1 255.255.255.0
ipv6 address 56:78:100::1/64
ipv6 address 56:78:102::1/64
mac-address 22.22.22
!

/* Configure CEF */ [ Required for Dual homing]

cef adjacency route override rib

/* Configure EVPN */

evpn
evi 102
bgp
rd 1605:102
route-target import 160102:102
route-target export 160102:102
!
advertise-mac
!
!
!
interface Bundle-Ether1
ethernet-segment
identifier type 0 56.56.56.56.56.56.56.01
!
!
interface Bundle-Ether2
ethernet-segment
identifier type 0 56.56.56.56.56.56.56.02
!
!

/* Configure L2VPN */

l2vpn
bridge group bg102
bridge-domain bd102
interface Bundle-Ether1.102
!
interface Bundle-Ether2.102
!
```



```

Originator: 160.0.0.6, Cluster list: 100.0.0.4
Path #2: Received by speaker 0
Advertised to update-groups (with more than one peer):
0.2
Local
56:78:100::2 from :: (160.0.0.5)
Origin incomplete, metric 0, localpref 100, weight 32768, valid, redistributed, best,
group-best
Received Path ID 0, Local Path ID 0, version 9
mac: 10:11:04:64:f2:7f

```

```
/* Verify Ethernet Segments are peering for Dual homing */
```

```
Router# show evpn ethernet-segment int bundle-Ether 1
```

```
Ethernet Segment Id Interface Nextthops
```

```
-----
0056.5656.5656.5656.5601 BE1 160.5.5.5
                               160.6.6.6
-----
```

```
/* Verify DF election */
```

```
Router# show evpn ethernet-segment int bundle-Ether 1 carving detail
```

```
Legend:
```

```

A - Load-balancing mode and Access Protection incompatible,
B - No Forwarders EVPN-enabled,
C - Backbone Source MAC missing (PBB-EVPN),
RT - ES-Import Route Target missing,
E - ESI missing,
H - Interface handle missing,
I - Name (Interface or Virtual Access) missing,
M - Interface in Down state,
O - BGP End of Download missing,
P - Interface already Access Protected,
Pf - Interface forced single-homed,
R - BGP RID not received,
S - Interface in redundancy standby state,
X - ESI-extracted MAC Conflict
SHG - No local split-horizon-group label allocated

```

```
Ethernet Segment Id Interface Nextthops
```

```
-----
0056.5656.5656.5656.5601 BE1 160.5.5.5
160.6.6.6
ES to BGP Gates : Ready
ES to L2FIB Gates : Ready
Main port :
Interface name : Bundle-Ether1
Interface MAC : 008a.9644.acdd
IfHandle : 0x080004dc
State : Up
Redundancy : Not Defined
ESI type : 0
Value : 56.5656.5656.5656.5601
ES Import RT : 5656.5656.5656 (from ESI)
Source MAC : 0000.0000.0000 (N/A)
Topology :
Operational : MH
Configured : All-active (AApF) (default)
Primary Services : Auto-selection
Secondary Services: Auto-selection
Service Carving Results:
Forwarders : 161

```

```
Permanent : 10
EVI:ETag P : 700:1, 701:1, 702:1, 703:1, 704:1, 705:1
EVI:ETag P : 706:1, 707:1, 708:1, 709:1
Elected : 76
EVI E : 100, 102, 104, 106, 108, 110
EVI E : 112, 114, 116, 118, 120, 122,
EVI E : 124, 126, 128, 130, 132, 134,
EVI E : 136, 138, 140, 142, 144, 146,
EVI E : 148, 150, 152, 154, 156, 158,
EVI E : 160, 162, 164, 166, 168, 170,
EVI E : 172, 174, 176, 178, 180, 182,
EVI E : 184, 186, 188, 190, 192, 194,
EVI E : 196, 198, 200, 202, 204, 206,
EVI E : 208, 210, 212, 214, 216, 218,
EVI E : 220, 222, 224, 226, 228, 230,
EVI E : 232, 234, 236, 238, 240, 242,
EVI E : 244, 246, 248, 250
Not Elected : 75
EVI NE : 101, 103, 105, 107, 109, 111
EVI NE : 113, 115, 117, 119, 121, 123,
EVI NE : 125, 127, 129, 131, 133, 135,
EVI NE : 137, 139, 141, 143, 145, 147,
EVI NE : 149, 151, 153, 155, 157, 159,
EVI NE : 161, 163, 165, 167, 169, 171,
EVI NE : 173, 175, 177, 179, 181, 183,
EVI NE : 185, 187, 189, 191, 193, 195,
EVI NE : 197, 199, 201, 203, 205, 207,
EVI NE : 209, 211, 213, 215, 217, 219,
EVI NE : 221, 223, 225, 227, 229, 231,
EVI NE : 233, 235, 237, 239, 241, 243,
EVI NE : 245, 247, 249
MAC Flushing mode : STP-TCN
Peering timer : 3 sec [not running]
Recovery timer : 30 sec [not running]
Carving timer : 0 sec [not running]
Local SHG label : 68663
Remote SHG labels : 1
68670 : nexthop 160.6.6.6
```




CHAPTER 9

EVPN Virtual Private Wire Service (VPWS)

The EVPN-VPWS is a BGP control plane solution for point-to-point services. It implements the signaling and encapsulation techniques for establishing an EVPN instance between a pair of PEs. It has the ability to forward traffic from one network to another without MAC lookup. The use of EVPN for VPWS eliminates the need for signaling single-segment and multi-segment PWs for point-to-point Ethernet services. The EVPN-VPWS technology works on IP and MPLS core; IP core to support BGP and MPLS core for switching packets between the endpoints.


EVPN-VPWS support both single-homing and multi-homing.

- [EVPN-VPWS Single Homed, on page 157](#)
- [EVPN-VPWS Multi-Homed, on page 159](#)

EVPN-VPWS Single Homed

The EVPN-VPWS single homed solution requires per EVI Ethernet Auto Discovery route. EVPN defines a new BGP Network Layer Reachability Information (NLRI) used to carry all EVPN routes. BGP Capabilities Advertisement used to ensure that two speakers support EVPN NLRI (AFI 25, SAFI 70) as per RFC 4760.

The architecture for EVPN VPWS is that the PEs run Multi-Protocol BGP in control-plane. The following image describes the EVPN-VPWS configuration:



- The VPWS service on PE1 requires the following three elements to be specified at configuration time:
 - The VPN ID (EVI)
 - The local AC identifier (AC1) that identifies the local end of the emulated service.
 - The remote AC identifier (AC2) that identifies the remote end of the emulated service.

PE1 allocates a MPLS label per local AC for reachability.

- The VPWS service on PE2 is set in the same manner as PE1. The three same elements are required and the service configuration must be symmetric.

PE2 allocates a MPLS label per local AC for reachability.

- PE1 advertises a single EVPN per EVI Ethernet AD route for each local endpoint (AC) to remote PEs with the associated MPLS label.

PE2 performs the same task.

- On reception of EVPN per EVI EAD route from PE2, PE1 adds the entry to its local L2 RIB. PE1 knows the path list to reach AC2, for example, next hop is PE2 IP address and MPLS label for AC2.

PE2 performs the same task.

Configure EVPN-VPWS Single Homed

This section describes how you can configure single-homed EVPN-VPWS feature.

```
Router# configure
Router(config)# router bgp 100
Router(config-bgp)# address-family l2vpn evpn
Router(config-bgp-af)# exit
Router(config-bgp-af)# neighbor 10.10.10.1
Router(config-bgp-af)# commit
Router(config-bgp-af)# exit
Router(config-bgp)# exit
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn-vpws
Router(config-l2vpn-xc)# p2p evpn1
Router(config-l2vpn-xc-p2p)# interface TenGigE0/1/0/2
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 100 target 12 source 10

Router(config-l2vpn-xc-p2p-pw)# commit
Router(config-l2vpn-xc-p2p)# exit
```

Running Configuration

```
configure
router bgp 100
  address-family l2vpn evpn
    neighbor 10.10.10.1
  !
!


configure
l2vpn
  xconnect group evpn-vpws
  p2p evpn1
    interface TenGigE0/1/0/2
      neighbor evpn evi 100 target 12 source 10
    !
  !
!
```


EVPN-VPWS Multi-Homed

The EVPN VPWS feature supports all-active multihoming capability that enables you to connect a customer edge device to two or more provider edge (PE) devices to provide load balancing and redundant connectivity. The load balancing is done using equal-cost multipath (ECMP).

When a CE device is multi-homed to two or more PEs and when all PEs can forward traffic to and from the multi-homed device for the VLAN, then such multihoming is referred to as all-active multihoming.

Figure 23: EVPN VPWS Multi-Homed



Consider the topology in which CE1 is multi-homed to PE1 and PE2; CE2 is multi-homed to PE3 and PE4. PE1 and PE2 will advertise an EAD per EVI route per AC to remote PEs which is PE3 and PE4, with the associated MPLS label. The ES-EAD route is advertised per ES (main interface), and it will not have a label. Similarly, PE3 and PE4 advertise an EAD per EVI route per AC to remote PEs, which is PE1 and PE2, with the associated MPLS label.

Consider a traffic flow from CE1 to CE2. Traffic is sent to either PE1 or PE2. The selection of path is dependent on the CE implementation for forwarding over a LAG. Traffic is encapsulated at each PE and forwarded to the remote PEs (PE 3 and PE4) through MPLS core. Selection of the destination PE is established by flow-based load balancing. PE3 and PE4 send the traffic to CE2. The selection of path from PE3 or PE4 to CE2 is established by flow-based load balancing.

If there is a failure and when the link from CE1 to PE1 goes down, the PE1 withdraws the ES-EAD route; sends a signal to the remote PEs to switch all the VPWS service instances associated with this multi-homed ES to backup PE, which is PE2.

Configure EVPN-VPWS Multi-Homed

This section describes how you can configure multi-homed EVPN-VPWS feature.

```
/* Configure PE1 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn_vpws
Router(config-l2vpn-xc)# p2p e1_5-6
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether10.2
```

```

Router(config-l2vpn-xc-p2p)# neighbor evpn evi 1 target 5 source 6

Router(config-l2vpn-xc-p2p-pw)# exit
Router(config-l2vpn-xc)# exit
Router(config-l2vpn)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether10
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.0a.00
Router(config-evpn-ac-es)# commit

/* Configure PE2 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn_vpws
Router(config-l2vpn-xc)# p2p e1_5-6
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether10.2
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 1 target 5 source 6

Router(config-l2vpn-xc-p2p-pw)# exit
Router(config-l2vpn-xc)# exit
Router(config-l2vpn)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether10
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.0a.00
Router(config-evpn-ac-es)# commit

/* Configure PE3 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn_vpws
Router(config-l2vpn-xc)# p2p e1_5-6
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether20.1
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 1 target 6 source 5

Router(config-l2vpn-xc-p2p-pw)# exit
Router(config-l2vpn-xc)# exit
Router(config-l2vpn)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether20
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.14.00
Router(config-evpn-ac-es)# commit

/* Configure PE4 */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn_vpws
Router(config-l2vpn-xc)# p2p e1_5-6
Router(config-l2vpn-xc-p2p)# interface Bundle-Ether20.1
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 1 target 6 source 5
Router(config-l2vpn-xc-p2p)# exit
Router(config-l2vpn-xc)# exit
Router(config-l2vpn)# exit
Router(config)# evpn
Router(config-evpn)# interface Bundle-Ether20
Router(config-evpn-ac)# ethernet-segment
Router(config-evpn-ac-es)# identifier type 0 00.01.00.ac.ce.55.00.14.00
Router(config-evpn-ac-es)# commit

```

Running Configuration

```
/* On PE1 */
!
configure
l2vpn xconnect group evpn_vpws
p2p e1_5-6
  interface Bundle-Ether10.2
  neighbor evpn evi 1 target 5 source 6
!
evpn
interface Bundle-Ether10
  ethernet-segment
  identifier type 0 00.01.00.ac.ce.55.00.0a.00
!

/* On PE2 */
!
configure
l2vpn xconnect group evpn_vpws
p2p e1_5-6
  interface Bundle-Ether10.2
  neighbor evpn evi 1 target 5 source 6
!
evpn
interface Bundle-Ether10
  ethernet-segment
  identifier type 0 00.01.00.ac.ce.55.00.0a.00
!

/* On PE3 */
!
configure
l2vpn xconnect group evpn_vpws
p2p e1_5-6
  interface Bundle-Ether20.1
  neighbor evpn evi 1 target 6 source 5
!
evpn
interface Bundle-Ether20
  ethernet-segment
  identifier type 0 00.01.00.ac.ce.55.00.14.00
!

/* On PE4 */
!
configure
l2vpn xconnect group evpn_vpws
p2p e1_5-6
  interface Bundle-Ether20.1
  neighbor evpn evi 1 target 6 source 5
!
evpn
interface Bundle-Ether20
  ethernet-segment
  identifier type 0 00.01.00.ac.ce.55.00.14.00
!
```




CHAPTER 10

L2VPN Services over Segment Routing for Traffic Engineering Policy

Segment Routing (SR) is a flexible and scalable way of performing source routing. The source device selects a path and encodes it in the packet header as an ordered list of segments. Segments are identifiers for any type of instruction.

Segment routing for traffic engineering (SR-TE) takes place through a tunnel between a source and destination pair. SR-TE uses the concept of source routing, where the source calculates the path and encodes it in the packet header as a segment. In SR-TE preferred path, each segment is an end-to-end path from the source to the destination, and instructs the routers in the provider core network to follow the specified path instead of the shortest path calculated by the IGP. The destination is unaware of the presence of the tunnel.

The user can achieve better resilience and convergence for the network traffic, by transporting MPLS L2VPN services using segment routing, instead of MPLS LDP. Segment routing can be directly applied to the MPLS architecture without changing the forwarding plane. In a segment-routing network that uses the MPLS data plane, LDP or other signaling protocol is not required; instead label distribution is performed by IGP. Removing protocols from the network simplifies its operation and makes it more robust and stable by eliminating the need for protocol interaction. Segment routing utilizes the network bandwidth more effectively than traditional MPLS networks and offers lower latency.

Preferred tunnel path functionality allows you map pseudowires to specific traffic-engineering tunnel paths. Attachment circuits are cross-connected to specific SR traffic engineering tunnel interfaces instead of remote PE router IP addresses reachable using IGP or LDP. Using preferred tunnel path, the traffic engineering tunnel transports traffic between the source and destination PE routers. A path is selected for an SR Policy when the path is valid and its preference is the best (highest value) among all the candidate paths of the SR Policy.

The following L2VPN services are supported over SR-TE policy:

- [EVPN VPWS Preferred Path over SR-TE Policy, on page 164](#)
- [L2VPN VPWS Preferred Path over SR-TE Policy, on page 176](#)
- [EVPN VPWS On-Demand Next Hop with SR-TE, on page 189](#)
- [Overview of Segment Routing , on page 203](#)
- [How Segment Routing Works , on page 204](#)
- [Segment Routing Global Block , on page 205](#)

EVPN VPWS Preferred Path over SR-TE Policy

EVPN VPWS Preferred Path over SR-TE Policy feature allows you to set the preferred path between the two end-points for EVPN VPWS pseudowire (PW) using SR-TE policy. SR policy allows you to choose the path on a per EVPN instance (EVI) basis. This feature is supported on bundle attachment circuit (AC) and physical AC.

Restrictions


- If EVPN VPWS with On Demand Next Hop (ODN) is configured, and EVPN VPWS with preferred path is also configured for the same PW, then the preferred-path will take precedence.
- EVPN VPWS SR policy is not supported on EVPN VPWS dual homing.
- EVPN validates if the route is for a single home next hop, otherwise it issues an error message about a dangling SR TE policy, and continue to set up EVPN-VPWS without it. EVPN relies on ESI value being zero to determine if this is a single home or not. If the AC is a Bundle-Ether interface running LACP then you need to manually configure the ESI value to zero to overwrite the auto-sense ESI as EVPN VPWS multihoming is not supported.

To disable EVPN dual homing, configure bundle-Ether AC with ESI value set to zero.

```
evpn
interface Bundle-Ether12
  ethernet-segment
    identifier type 0 00.00.00.00.00.00.00.00
/* Or globally */
Evpn
  ethernet-segment type 1 auto-generation-disable
```

Topology

Figure 24: EVPN VPWS Preferred Path over SR-TE Policy



Consider a topology where PE1 and PE3 are the two EVPN VPWS PW end-points. Traffic is sent from PE1 to PE3 through SR in the core. Traffic from PE1 can be sent to PE3 either through P1 or P2 node. In this example, the EVPN VPWS preferred path over SR policy is configured to show the traffic flow from PE1 to PE3 using prefix-SID. Using adjacency-SID, you can steer traffic flow from PE1 to PE3 and specify whether it should pass through P1 or P2 node.

Configure EVPN VPWS Preferred Path over SR-TE Policy

You must complete these tasks to ensure the successful configuration of EVPN VPWS Preferred Path over SR-TE Policy feature:

- Configure Prefix-SID on IGP — The following examples show how to configure prefix-SID in IS-IS.
- Configure Adjacency-SID on IGP — The following examples show how to configure Adjacency-SID in IS-IS.
- Configure segment-list
- Configure SR-TE policy
- Configure EVPN VPWS over SR-TE policy

Configure Prefix-SID in ISIS

Configure Prefix-SID on PE1, P1, P2, and PE3.

```

/* Configure Prefix-SID on PE1 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 180000 200000
Router(config-sr)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# is-type level-2-only
Route(config-isis)# net 49.0002.0330.2000.0031.00
Route(config-isis)# nsr
Route(config-isis)# nsf ietf
Route(config-isis)# log adjacency changes
Route(config-isis)# address-family ipv4 unicast
Route(config-isis-af)# metric-style wide level 2
Route(config-isis-af)# mpls traffic-eng level-2-only
Route(config-isis-af)# mpls traffic-eng router-id 1.1.1.1
Route(config-isis-af)# segment-routing mpls sr-prefer
Route(config-isis-af)# segment-routing prefix-sid-map advertise-local
Route(config-isis-af)# exit
!
Route(config-isis)# interface loopback 0
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-af)# prefix-sid index 180010
Route(config-isis-af)# commit
Route(config-isis-af)# exit

/* Configure Prefix-SID on P1 */

Router# configure
Router(config)# segment-routing

```

```

Router(config-sr)# global-block 180000 200000
Router(config-sr)# exit
!
Router# configure
Router(config)# router isis core
Router(config-isis)# is-type level-2-only
Router(config-isis)# net 49.0002.0330.2000.0021.00
Router(config-isis)# nsr
Router(config-isis)# nsf ietf
Router(config-isis)# log adjacency changes
Router(config-isis)# address-family ipv4 unicast
Router(config-isis-af)# metric-style wide level 2
Router(config-isis-af)# mpls traffic-eng level-2-only
Router(config-isis-af)# mpls traffic-eng router-id loopback0
Router(config-isis-af)# segment-routing mpls sr-prefer
Router(config-isis-af)# segment-routing prefix-sid-map advertise-local
Router(config-isis-af)# exit
!
Router(config-isis)# interface loopback 0
Router(config-isis-if)# address-family ipv4 unicast
Router(config-isis-af)# prefix-sid index 180015
Router(config-isis-af)# commit
Router(config-isis-af)# exit

/* Configure Prefix-SID on P2 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 180000 200000
Router(config-sr)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# is-type level-2-only
Route(config-isis)# net 49.0002.0330.2000.0022.00
Route(config-isis)# nsr
Route(config-isis)# nsf ietf
Route(config-isis)# log adjacency changes
Route(config-isis)# address-family ipv4 unicast
Route(config-isis-af)# metric-style wide level 2
Route(config-isis-af)# mpls traffic-eng level-2-only
Route(config-isis-af)# mpls traffic-eng router-id loopback0
Route(config-isis-af)# segment-routing mpls sr-prefer
Route(config-isis-af)# segment-routing prefix-sid-map advertise-local
Route(config-isis-af)# exit
!
Route(config-isis)# interface loopback 0
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-af)# prefix-sid index 180025
Route(config-isis-af)# commit
Route(config-isis-af)# exit

/* Configure Prefix-SID on PE3 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 180000 200000
Router(config-sr)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# is-type level-2-only
Route(config-isis)# net 49.0002.0330.2000.3030.0030.0035.00

```



```

Route(config-isis)# address-family ipv4 unicast
Route(config-isis-af)# metric-style wide level 2
Route(config-isis-af)# mpls traffic-eng level-2-only
Route(config-isis-af)# mpls traffic-eng router-id loopback0
Route(config-isis-af)# segment-routing mpls sr-prefer
Route(config-isis-af)# segment-routing prefix-sid-map advertise-local
Route(config-isis-af)# exit
!
Route(config-isis)# interface loopback0
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-af)# prefix-sid index 180020
Route(config-isis-af)# commit
Route(config-isis-af)# exit

```

Configure Adjacency-SID in ISIS

Configure Adjacency-SID on PE1, P1, P2, and PE3.

```

/* Configure Adjacency-SID on PE1 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# local-block 15000 15999
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface Bundle-Ether121
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15101
Route(config-isis-if-af)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/1/6
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15102
Route(config-isis-if-af)# commit

/* Configure Adjacency-SID on P1 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# local-block 15000 15999
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface Bundle-Ether121
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# metric 20
Route(config-isis-if-af)# adjacency-sid absolute 15200
Route(config-isis-if-af)# commit

```

```

!
Router# configure
Router(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/7
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15202
Route(config-isis-if-af)# commit
!
/* Configure Adjacency-SID on P2 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# local-block 15000 15999
!
Router# configure
Router(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/7
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# metric 20
Route(config-isis-if-af)# adjacency-sid absolute 15201
Route(config-isis-if-af)# exit
!
Router# configure
Router(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/5
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# metric 20
Route(config-isis-if-af)# adjacency-sid absolute 15204
Route(config-isis-if-af)# commit

/* Configure Adjacency-SID on PE3 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# local-block 15000 15999
!
Router# configure
Router(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/1
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15301
Route(config-isis-if-af)# exit
!
Router# configure
Router(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/2
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast

```

```
Route(config-isis-if-af)# adjacency-sid absolute 15302
Route(config-isis-if-af)# commit
```

Configure Segment-list

```
/* Configure Segment-list on PE1 using prefix-SID */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 180000 200000
Router(config-sr)# traffic-eng
Router(config-sr-te)# logging
Router(config-sr-te-log)# policy status
Router(config-sr-te-log)# exit
!
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list name pref_sid_to_PE3
Router(config-sr-te-sl)# index 1 mpls label 180020 <-----using prefix-SID
Router(config-sr-te-sl)# exit

/* Configure Segment-list on PE1 using adjacency-SID */

Router# configure
Router(config)# segment-routing
Router(config-sr)# local-block 15000 15999
Router(config-sr)# traffic-eng
Router(config-sr-te)# logging
Router(config-sr-te-log)# policy status
Router(config-sr-te-log)# exit
!
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list name pref_adj_sid_to_PE3
Router(config-sr-te-sl)# index 1 mpls label 15101 <-----using adjacency-SID
Router(config-sr-te-sl)# index 2 mpls label 15202 <-----using adjacency-SID
Router(config-sr-te-sl)# exit
```

Configure SR-TE Policy

```
/* Configure SR-TE Policy */

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy pref_sid_to_PE3
Router(config-sr-te-policy)# color 9001 end-point ipv4 20.20.20.20
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 10
Router(config-sr-te-policy-path-pref)# explicit segment-list pref_sid_to_PE3
Router(config-sr-te-policy-path-pref)# commit
Router(config-sr-te-pp-info)# exit
!
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy pref_adj_sid_to_PE3
Router(config-sr-te-policy)# color 9001 end-point ipv4 20.20.20.20
```

```

Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 200
Router(config-sr-te-policy-path-pref)# explicit segment-list pref_adj_sid_to_PE3
Router(config-sr-te-policy-path-pref)# commit
Router(config-sr-te-pp-info)# exit

/* You can configure multiple preferences for an SR policy. Among the configured preferences,
the largest number takes the highest precedence */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 180000 200000
Router(config-sr)# local-block 15000 15999
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy 1013
Router(config-sr-te-policy)# color 1013 end-point ipv4 2.2.2.2
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy-path)# preference 100
Router(config-sr-te-policy-path-pref)# explicit segment-list PE1-P1_BE121
Router(config-sr-te-policy-path-pref)# exit
!
Router(config-sr-te-policy-path)# preference 200
Router(config-sr-te-policy-path-pref)# explicit segment-list PE1-PE3-P1-t0016
Router(config-sr-te-policy-path-pref)# exit
!
Router(config-sr-te-policy-path)# preference 700 <-----largest number takes the
precedence
Router(config-sr-te-policy-path-pref)# explicit segment-list PE1-P1
Router(config-sr-te-policy-path-pref)# commit
Router(config-sr-te-policy-path-pref)# exit

```

Configure EVPN VPWS over SR-TE Policy

```

Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# pw-class 1001
Router(config-l2vpn-pwc)# encapsulation mpls
Router(config-l2vpn-pwc-mpls)# preferred-path sr-te policy pref_sid_to_PE3 fallback disable
Router(config-l2vpn-pwc-mpls)# commit
Router(config-l2vpn-pwc-mpls)# exit
!
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn_vpws
Router(config-l2vpn-xc)# p2p evpn_vpws_1001
Router(config-l2vpn-xc-p2p)# interface tengi0/1/0/1.1001
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 1001 target 10001 source 20001
Router(config-l2vpn-xc-p2p-pw)# pw-class 1001
Router(config-l2vpn-xc-p2p-pw)# commit
Router(config-l2vpn-xc-p2p-pw)# exit

/* If Fallback Enable is configured, which is the default option, and if the SR-policy is
down, then EVPN VPWS will still continue to be UP using the regular IGP path, and not using
the SR-policy */
show l2vpn xconnect detail
EVPN: neighbor 20.20.20.20, PW ID: evi 1001, ac-id 10001, state is up ( established )
Preferred path Inactive : SR TE pref_sid_to_PE3, Statically configured, fallback enabled

Tunnel : Down

```

LSP: Up

```
/* If Fallback Disable is configured, and if the SR-policy is down, or if it misconfigured
in dual homed mode, then the L2VPN PW will be down */
```

```
show l2vpn xconnect detail
```

```
EVPN: neighbor 20.20.20.20, PW ID: evi 1001, ac-id 10001, state is down ( local ready )
```

```
Preferred path Active : SR TE pref_sid_to_PE3, Statically configured, fallback disabled
```

```
Tunnel : Down
```

Running Configuration

```
/* Configure Prefix-SID in ISIS */
```

PE1:

```
configure
  segment-routing
    global-block 180000 200000
  !
router isis core
  is-type level-2-only
  net 49.0002.0330.2000.0031.00
  nsr
  nsf ietf
  log adjacency changes
  address-family ipv4 unicast
  metric-style wide level 2
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id 1.1.1.1
  segment-routing mpls sr-prefer
  segment-routing prefix-sid-map advertise-local

interface Loopback0
  address-family ipv4 unicast
  prefix-sid index 180010
```

P1:

```
configure
  segment-routing
    global-block 180000 200000

router isis core
  is-type level-2-only
  net 49.0002.0330.2000.0021.00
  nsr
  nsf ietf
  log adjacency changes
  address-family ipv4 unicast
  metric-style wide level 2
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id Loopback0
  segment-routing mpls sr-prefer
  segment-routing prefix-sid-map advertise-local

interface Loopback0
  address-family ipv4 unicast
  prefix-sid index 180015
```

P2:

```
configure
```

```

segment-routing
  global-block 180000 200000

router isis core
  is-type level-2-only
  net 49.0002.0330.2000.0022.00
  nsr
  nsf ietf
  log adjacency changes
  address-family ipv4 unicast
  metric-style wide level 2
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id Loopback0
  segment-routing mpls sr-prefer
  segment-routing prefix-sid-map advertise-local

interface Loopback0
  address-family ipv4 unicast
  prefix-sid index 180025

```

PE3:

```

configure
  segment-routing
    global-block 180000 200000

router isis core
  is-type level-2-only
  net 49.0002.0330.2000.3030.0035.00
  address-family ipv4 unicast
  metric-style wide level 2
  mpls traffic-eng level-2-only
  mpls traffic-eng router-id Loopback0
  segment-routing mpls sr-prefer
  segment-routing prefix-sid-map advertise-local

interface Loopback0
  address-family ipv4 unicast
  prefix-sid index 180020

/* Configure Adjacency-SID in ISIS */

```

PE1:

```

configure
  segment-routing
    local-block 15000 15999
  !

router isis core
  !
interface Bundle-Ether121
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15101

interface TenGigE0/0/1/6
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15102

```

PE1:

```
configure
  segment-routing
    local-block 15000 15999

router isis core
!
interface Bundle-Ether121
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  metric 20
  adjacency-sid absolute 15200

interface TenGigE0/0/0/0/7
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  metric 20
  adjacency-sid absolute 15202
```

PE2:

```
configure
  segment-routing
    local-block 15000 15999

router isis core
!
interface TenGigE0/0/0/0/5
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  metric 20
  adjacency-sid absolute 15204

interface TenGigE0/0/0/0/7
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  metric 20
  adjacency-sid absolute 15201
```

PE3:

```
configure
  segment-routing
    local-block 15000 15999

router isis core
!
interface TenGigE0/0/0/0/1
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15301
!
```

```

!
interface TenGigE0/0/0/2
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
    adjacency-sid absolute 15302

/* Configure Segment-list */
PE1:

configure
  segment-routing
    global-block 180000 200000
    traffic-eng
      logging
      policy status

segment-routing
  traffic-eng
    segment-list name pref_sid_to_PE3
      index 1 mpls label 180020
    !
  !
configure
  segment-routing
    local-block 15000 15999
    traffic-eng
      logging
      policy status

segment-routing
  traffic-eng
    segment-list name pref_adj_sid_to_PE3
      index 1 mpls label 15101
      index 2 mpls label 15202
    !
  !

/* Configure SR-TE policy */

segment-routing
  traffic-eng
    policy pref_sid_to_PE3
      color 9001 end-point ipv4 20.20.20.20
      candidate-paths
        preference 10
        explicit segment-list pref_sid_to_PE3
      !
    !
segment-routing
  traffic-eng
    policy pref_adj_sid_to_PE3
      color 9001 end-point ipv4 20.20.20.20
      candidate-paths
        preference 200
        explicit segment-list pref_adj_sid_to_PE3
      !
    !

/* You can configure multiple preferences for an SR policy. Among the configured preferences,
the largest number takes the highest precedence */

segment-routing

```



```

traffic-eng
policy 1013
color 1013 end-point ipv4 2.2.2.2
candidate-paths
preference 100
  explicit segment-list PE1-P1_BE121
  !
preference 200
  explicit segment-list PE1-PE3-P1-t0016
  !
preference 700
  explicit segment-list PE1-P1
  !

/* Configure EVPN VPWS over SR-TE policy */
PE1:
configure
l2vpn
pw-class 1001
  encapsulation mpls
  preferred-path sr-te policy pref_sid_to_PE3 fallback disable
xconnect group evpn_vpws
p2p evpn_vpws_1001
interface tengi0/1/0/1.1001
neighbor evpn evi 1001 target 10001 source 20001
pw-class 1001
!

```

Verify EVPN VPWS Preferred Path over SR-TE Policy Configuration

```

PE1#show segment-routing traffic-eng forwarding policy name pref_sid_to_PE3 detail
Policy          Segment          Outgoing          Outgoing          Next Hop          Bytes
Name            List             Label             Interface          Switched
-----
pref_sid_to_PE3
                                15102             TenGigE0/0/1/6    20.20.20.20      81950960
                                Label Stack (Top -> Bottom): { 15101, 15102 }
                                Path-id: 1, Weight: 0
                                Packets Switched: 787990
Local label: 34555
Packets/Bytes Switched: 1016545/105720680
(!): FRR pure backup

```

```

PE1#show mpls forwarding tunnels sr-policy name pref_sid_to_PE3
Tunnel          Outgoing          Outgoing          Next Hop          Bytes
Name            Label             Interface          Switched
-----
pref_sid_to_PE3 (SR) 15102 TenGigE0/0/1/6 20.20.20.20      836516512

```

```

PE1#show l2vpn xconnect group evpn_vpws xc-name evpn_vpws_1001 detail
Group evpn_vpws, XC evpn_vpws_1001, state is up; Interworking none
AC: Bundle-Ether12.1001, state is up
Type VLAN; Num Ranges: 1
Outer Tag: 1000
Rewrite Tags: []
VLAN ranges: [1, 1]
MTU 1500; XC ID 0xc0000018; interworking none
Statistics:
  packets: received 642304, sent 642244

```

```

bytes: received 61661184, sent 61655424
drops: illegal VLAN 0, illegal length 0
EVPN: neighbor 20.20.20.20, PW ID: evi 1001, ac-id 10001, state is up ( established )
XC ID 0xa0000007
Encapsulation MPLS
Source address 10.10.10.10
Encap type Ethernet, control word enabled
Sequencing not set
Preferred path Active : SR TE pref_sid_to_PE3, Statically configured, fallback disabled
Tunnel : Up
Load Balance Hashing: src-dst-mac

```

Associated Commands

- [adjacency-sid](#)
- [index](#)
- [prefix-sid](#)
- [router isis](#)
- [segment-routing](#)

The applicable segment routing commands are described in the *Segment Routing Command Reference for Cisco NCS 5500 Series Routers and Cisco NCS 540 Series Routers*

Related Topics

- [Overview of Segment Routing , on page 203](#)
- [How Segment Routing Works , on page 204](#)
- [Segment Routing Global Block , on page 205](#)

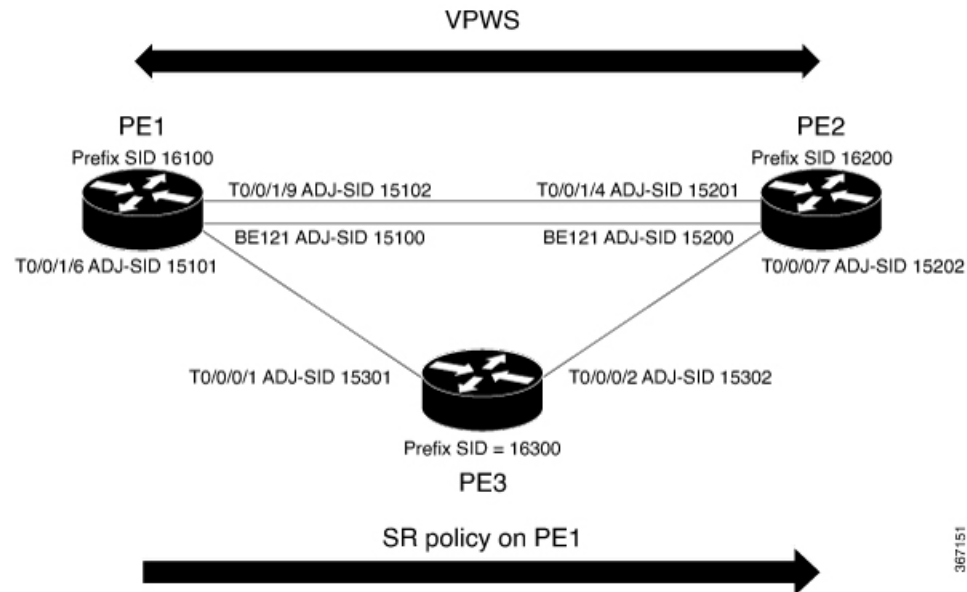
L2VPN VPWS Preferred Path over SR-TE Policy

L2VPN VPWS Preferred Path over SR-TE Policy feature allows you to set the preferred path between the two end-points for L2VPN Virtual Private Wire Service (VPWS) using SR-TE policy.

Configure L2VPN VPWS Preferred Path over SR-TE Policy

Perform the following steps to configure L2VPN VPWS Preferred Path over SR-TE Policy feature. The following figure is used as a reference to explain the configuration steps.

Figure 25: L2VPN VPWS Preferred Path over SR-TE Policy



- Configure Prefix-SID on IGP — The following examples show how to configure prefix-SID in IS-IS.
- Configure Adjacency-SID on IGP — The following examples show how to configure Adjacency-SID in IS-IS.
- Configure segment-list
- Configure SR-TE policy
- Configure VPWS over SR-TE policy

Configure Prefix-SID in IS-IS

Configure Prefix-SID on PE1, PE2, and PE3.

```

/* Configure Prefix-SID on PE1 */

Router# configure
Route(config)# router isis core
Route(config-isis)# is-type level-2-only
Route(config-isis)# net 49.0002.0330.2000.0031.00
Route(config-isis)# nsr
Route(config-isis)# nsf ietf
Route(config-isis)# log adjacency changes
Route(config-isis)# address-family ipv4 unicast
Route(config-isis-af)# metric-style wide level 2
Route(config-isis-af)# mpls traffic-eng level-2-only
Route(config-isis-af)# mpls traffic-eng router-id 1.1.1.1
Route(config-isis-af)# segment-routing mpls sr-prefer
Route(config-isis-af)# segment-routing prefix-sid-map advertise-local
Route(config-isis-af)# exit
!
Route(config-isis)# interface loopback 0
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-af)# prefix-sid index 16100

```

```

Route(config-isis-af) # commit
Route(config-isis-af) # exit

/* Configure Prefix-SID on PE2 */

Router# configure
Route(config) # router isis core
Route(config-isis) # is-type level-2-only
Route(config-isis) # net 49.0002.0330.2000.0021.00
Route(config-isis) # nsr
Route(config-isis) # nsf ietf
Route(config-isis) # log adjacency changes
Route(config-isis) # address-family ipv4 unicast
Route(config-isis-af) # metric-style wide level 2
Route(config-isis-af) # mpls traffic-eng level-2-only
Route(config-isis-af) # mpls traffic-eng router-id loopback0
Route(config-isis-af) # segment-routing mpls sr-prefer
Route(config-isis-af) # segment-routing prefix-sid-map advertise-local
Route(config-isis-af) # exit
!
Route(config-isis) # interface loopback 0
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-af) # prefix-sid index 16200
Route(config-isis-af) # commit
Route(config-isis-af) # exit

/* Configure Prefix-SID on PE3 */

Router# configure
Route(config) # router isis core
Route(config-isis) # is-type level-2-only
Route(config-isis) # net 49.0002.0330.2000.3030.0030.0035.00
Route(config-isis) # address-family ipv4 unicast
Route(config-isis-af) # metric-style wide level 2
Route(config-isis-af) # mpls traffic-eng level-2-only
Route(config-isis-af) # mpls traffic-eng router-id loopback0
Route(config-isis-af) # segment-routing mpls sr-prefer
Route(config-isis-af) # segment-routing prefix-sid-map advertise-local
Route(config-isis-af) # exit
!
Route(config-isis) # interface loopback 0
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-af) # prefix-sid index 16300
Route(config-isis-af) # commit
Route(config-isis-af) # exit

```

Configure Adjacency-SID in IS-IS

Configure Adjacency-SID on PE1, PE2, and PE3.

```

/* Configure Adjacency-SID on PE1 */

Router# configure
Route(config) # router isis core
Route(config-isis) # interface Bundle-Ether121
Route(config-isis-if) # circuit-type level-2-only
Route(config-isis-if) # point-to-point
Route(config-isis-if) # hello-padding disable
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-if-af) # adjacency-sid absolute 15100

```

```
Route(config-isis-if-af)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/1
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15101
Route(config-isis-if-af)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/2
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15102
Route(config-isis-if-af)# commit

/* Configure Adjacency-SID on PE2 */

Router# configure
Route(config)# router isis core
Route(config-isis)# interface Bundle-Ether121
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15200
Route(config-isis-if-af)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/3
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15201
Route(config-isis-if-af)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/7
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15202
Route(config-isis-if-af)# commit

/* Configure Adjacency-SID on PE3 */

Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/1
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
```

```

Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15301
Route(config-isis-if-af)# exit
!
Router# configure
Route(config)# router isis core
Route(config-isis)# interface TenGigE0/0/0/2
Route(config-isis-if)# circuit-type level-2-only
Route(config-isis-if)# point-to-point
Route(config-isis-if)# hello-padding disable
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-if-af)# adjacency-sid absolute 15302
Route(config-isis-if-af)# commit

```

Configure Segment-list

Configure segment-list on PE1, PE2, and PE3.

```

/* Configure segment-list on PE1 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 16000 23999
Router(config-sr)# local-block 15000 15999
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list segment-list name PE1-PE2
Router(config-sr-te-sl)# index 1 mpls label 16200
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE3
Router(config-sr-te-sl)# index 1 mpls label 16300
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE2-PE3
Router(config-sr-te-sl)# index 1 mpls label 16200
Router(config-sr-te-sl)# index 2 mpls label 16300
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE2_bad
Router(config-sr-te-sl)# index 1 mpls label 16900
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE3-PE2
Router(config-sr-te-sl)# index 1 mpls label 16300
Router(config-sr-te-sl)# index 2 mpls label 16200
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE2_BE121
Router(config-sr-te-sl)# index 1 mpls label 15100
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE3-PE2_link
Router(config-sr-te-sl)# index 1 mpls label 15101
Router(config-sr-te-sl)# index 2 mpls label 15302
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE1-PE3-PE2-t0016
Router(config-sr-te-sl)# index 1 mpls label 15101
Router(config-sr-te-sl)# index 2 mpls label 16200
Router(config-sr-te-sl)# commit

```

```

/* Configure segment-list on PE2 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 16000 23999
Router(config-sr)# local-block 15000 15999
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list segment-list name PE2-PE1
Router(config-sr-te-sl)# index 1 mpls label 16100
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE2-PE3-PE1
Router(config-sr-te-sl)# index 1 mpls label 16300
Router(config-sr-te-sl)# index 2 mpls label 16100
Router(config-sr-te-sl)# commit

/* Configure segment-list on PE3 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# global-block 16000 23999
Router(config-sr)# local-block 15000 15999
Router(config-sr)# traffic-eng
Router(config-sr-te)# segment-list segment-list name PE3-PE1
Router(config-sr-te-sl)# index 1 mpls label 16100
Router(config-sr-te-sl)# exit
!
Router(config-sr-te)# segment-list segment-list name PE3-PE2-PE1
Router(config-sr-te-sl)# index 1 mpls label 16200
Router(config-sr-te-sl)# index 2 mpls label 16100
Router(config-sr-te-sl)# commit

```

Configure SR-TE Policy

```

/* Configure SR-TE policy */

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy 100
Router(config-sr-te-policy)# color 1 end-point ipv4 2.2.2.2
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy)# preference 400
Router(config-sr-te-pp-info)# explicit segment-list PE1-PE3-PE2
Router(config-sr-te-pp-info)# exit
!
Router(config-sr-te-policy)# preference 500 <-----largest number takes the
precedence
Router(config-sr-te-pp-info)# explicit segment-list PE1-PE2
Router(config-sr-te-pp-info)# commit
Router(config-sr-te-pp-info)# exit
!
Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# policy 1013
Router(config-sr-te-policy)# color 1013 end-point ipv4 2.2.2.2
Router(config-sr-te-policy)# candidate-paths
Router(config-sr-te-policy)# preference 100

```

```

Router(config-sr-te-pp-info) # explicit segment-list PE1-PE2_BE121
Router(config-sr-te-pp-info) # exit
!
Router(config-sr-te-policy) # preference 200
Router(config-sr-te-pp-info) # explicit segment-list PE1-PE3-PE2-t0016
Router(config-sr-te-pp-info) # exit
!
Router(config-sr-te-policy) # preference 500
Router(config-sr-te-pp-info) # explicit segment-list PE1-PE2
Router(config-sr-te-pp-info) # exit
!
Router(config-sr-te-policy) # preference 600
Router(config-sr-te-pp-info) # explicit segment-list PE1-PE3-PE2
Router(config-sr-te-pp-info) # exit
!
Router(config-sr-te-policy) # preference 700
Router(config-sr-te-pp-info) # explicit segment-list PE1-PE3-PE2_link
Router(config-sr-te-pp-info) # commit
!
Router# configure
Router(config) # segment-routing
Router(config-sr) # traffic-eng
Router(config-sr-te) # policy 1300
Router(config-sr-te-policy) # color 1300 end-point ipv4 3.3.3.3
Router(config-sr-te-policy) # candidate-paths
Router(config-sr-te-policy) # preference 100
Router(config-sr-te-pp-info) # explicit segment-list PE1-PE3
Router(config-sr-te-pp-info) # commit
!

```

Configure VPWS over SR-TE Policy

```

Router# configure
Router(config) # l2vpn
Router(config-l2vpn) # pw-class pw1300
Router(config-l2vpn-pwc) # encapsulation mpls
Router(config-l2vpn-pwc-mpls) # load-balancing
Router(config-l2vpn-pwc-mpls-load-bal) # flow-label both
Router(config-l2vpn-pwc-mpls-load-bal) # exit
!
Router(config-l2vpn-pwc-mpls) # preferred-path sr-te policy 1300 fallback disable
Router(config-l2vpn-pwc-mpls) # exit
!
Router(config) # l2vpn
Router(config-l2vpn) # xconnect group xcon1
Router(config-l2vpn-xc) # p2p vplw1002
Router(config-l2vpn-xc-p2p) # interface TenGigE0/0/0/5
Router(config-l2vpn-xc-p2p) # neighbor 3.3.3.3 pw-id 1002
Router(config-l2vpn-xc-p2p-pw) # pw-class pw1300
Router(config-l2vpn-xc-p2p-pw) # commit
Router(config-l2vpn-xc-p2p-pw) # exit

```

Running Configuration

```

/* Configure prefix-SID */
PE1:
router isis core
 is-type level-2-only
 net 49.0002.0330.2000.0031.00

```



```

nsr
nsf ietf
log adjacency changes
address-family ipv4 unicast
metric-style wide level 2
mpls traffic-eng level-2-only
mpls traffic-eng router-id 1.1.1.1
segment-routing mpls sr-prefer
segment-routing prefix-sid-map advertise-local

interface Loopback0
address-family ipv4 unicast
prefix-sid index 16100

```

PE2:

```

router isis core
is-type level-2-only
net 49.0002.0330.2000.0021.00
nsr
nsf ietf
log adjacency changes
address-family ipv4 unicast
metric-style wide level 2
mpls traffic-eng level-2-only
mpls traffic-eng router-id Loopback0
segment-routing mpls sr-prefer
segment-routing prefix-sid-map advertise-local

interface Loopback0
address-family ipv4 unicast
prefix-sid index 16200

```

PE3:

```

router isis core
is-type level-2-only
net 49.0002.0330.2000.3030.0030.0035.00
address-family ipv4 unicast
metric-style wide level 2
mpls traffic-eng level-2-only
mpls traffic-eng router-id Loopback0
segment-routing mpls sr-prefer
segment-routing prefix-sid-map advertise-local

interface Loopback0
address-family ipv4 unicast
prefix-sid index 16300

```

```
/* Configure Adjacency-SID */
```

PE1:

```

router isis core
!
interface Bundle-Ether121
circuit-type level-2-only
point-to-point
hello-padding disable
address-family ipv4 unicast
adjacency-sid absolute 15100
!
interface TenGigE0/0/0/1

circuit-type level-2-only
point-to-point
hello-padding disable
address-family ipv4 unicast

```

```

    adjacency-sid absolute 15101
    !
interface TenGigE0/0/0/2
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15102

PE2
router isis core
!
interface Bundle-Ether121
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15200

interface TenGigE0/0/0/0/4
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15201

interface TenGigE0/0/0/0/7
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15202

PE3:
router isis core
!
interface TenGigE0/0/0/1
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15301
!
!
interface TenGigE0/0/0/2
  circuit-type level-2-only
  point-to-point
  hello-padding disable
  address-family ipv4 unicast
  adjacency-sid absolute 15302

/* Configure segment-list */
PE1:
segment-routing
global-block 16000 23999
local-block 15000 15999
traffic-eng
segment-list name PE1-PE2
  index 1 mpls label 16200
!
segment-list name PE1-PE3
  index 1 mpls label 16300
!
segment-list name PE1-PE2-PE3

```

```

    index 1 mpls label 16200
    index 2 mpls label 16300
    !
segment-list name PE1-PE2_bad
    index 1 mpls label 16900
    !
segment-list name PE1-PE3-PE2
    index 1 mpls label 16300
    index 2 mpls label 16200
    !
segment-list name PE1-PE2_BE121
    index 1 mpls label 15100
!
segment-list name PE1-PE3-PE2_link
    index 1 mpls label 15101
    index 2 mpls label 15302
    !

segment-list name PE1-PE3-PE2-t0016
    index 1 mpls label 15101
    index 2 mpls label 16200

PE2:
segment-routing
global-block 16000 23999
local-block 15000 15999
traffic-eng
segment-list name PE2-PE1
    index 1 mpls label 16100
    !
segment-list name PE2-PE3-PE1
    index 1 mpls label 16300
    index 2 mpls label 16100

PE3:
segment-routing
global-block 16000 23999
local-block 15000 15999
traffic-eng
segment-list name PE3-PE1
    index 1 mpls label 16100
    !
segment-list name PE3-PE2-PE1
    index 1 mpls label 16200
    index 2 mpls label 16100

/* Configure SR-TE policy */

segment-routing
traffic-eng
policy 100
    color 1 end-point ipv4 2.2.2.2
    candidate-paths
    preference 400
        explicit segment-list PE1-PE3-PE2
        !
    preference 500
        explicit segment-list PE1-PE2

policy 1013
    color 1013 end-point ipv4 2.2.2.2
    candidate-paths
    preference 100
        explicit segment-list PE1-PE2_BE121

```

Verify L2VPN VPWS Preferred Path over SR-TE Policy Configuration

```

!
preference 200
  explicit segment-list PE1-PE3-PE2-t0016
!
preference 500
  explicit segment-list PE1-PE2
!
preference 600
  explicit segment-list PE1-PE3-PE2
!
preference 700
  explicit segment-list PE1-PE3-PE2_link
!
policy 1300
  color 1300 end-point ipv4 3.3.3.3
  candidate-paths
  preference 100
  explicit segment-list PE1-PE3
!

/*Configure VPWS over SR-TE policy
l2vpn
pw-class pw1300
  encapsulation mpls
  load-balancing
  flow-label both
  preferred-path sr-te policy 1300 fallback disable

Xconnect group xcon1
  p2p vplw1002
  interface TenGigE0/0/0/5
  neighbor 3.3.3.3 pw-id 1002
  pw-class pw1300

```

Verify L2VPN VPWS Preferred Path over SR-TE Policy Configuration

```

/* The prefix-sid and Adjacency-sid must be in the SR topology */

PE1#show segment-routing traffic-eng ipv4 topology | inc Prefix
Thu Feb  1 20:28:43.343 EST
Prefix SID:
  Prefix 1.1.1.1, label 16100 (regular)
Prefix SID:
  Prefix 3.3.3.3, label 16300 (regular)
Prefix SID:
  Prefix 2.2.2.2, label 16200 (regular)

PE1#show segment-routing traffic-eng ipv4 topology | inc Adj SID
Thu Feb  1 20:30:25.760 EST
Adj SID: 61025 (unprotected) 15102 (unprotected)
Adj SID: 61023 (unprotected) 15101 (unprotected)
Adj SID: 65051 (unprotected) 15100 (unprotected)
Adj SID: 41516 (unprotected) 15301 (unprotected)
Adj SID: 41519 (unprotected) 15302 (unprotected)
Adj SID: 46660 (unprotected) 15201 (unprotected)
Adj SID: 24003 (unprotected) 15202 (unprotected)
Adj SID: 46675 (unprotected) 15200 (unprotected)
PE1# show segment-routing traffic-eng policy candidate-path name 100

```

```
SR-TE policy database
-----
```

```
Color: 100, End-point: 2.2.2.2
Name: srte_c_1_ep_2.2.2.2
```

```
PE1#show segment-routing traffic-eng policy name 100
Thu Feb 1 23:16:58.368 EST
```

```
SR-TE policy database
-----
```

```
Name: 100 (Color: 1, End-point: 2.2.2.2)
Status:
  Admin: up Operational: up for 05:44:25 (since Feb 1 17:32:34.434)
Candidate-paths:
  Preference 500:
    Explicit: segment-list PE1-PE2 (active)
      Weight: 0, Metric Type: IGP
      16200 [Prefix-SID, 2.2.2.2]
  Preference 400:
    Explicit: segment-list PE1-PE3-PE2 (inactive)
    Inactive Reason: unresolved first label
    Weight: 0, Metric Type: IGP
Attributes:
  Binding SID: 27498
  Allocation mode: dynamic
  State: Programmed
  Policy selected: yes
  Forward Class: 0
```

```
PE1#show segment-routing traffic-eng policy name 1013
Thu Feb 1 21:20:57.439 EST
```

```
SR-TE policy database
-----
```

```
Name: 1013 (Color: 1013, End-point: 2.2.2.2)
Status:
  Admin: up Operational: up for 00:06:36 (since Feb 1 21:14:22.057)
Candidate-paths:
  Preference 700:
    Explicit: segment-list PE1-PE3-PE2_link (active)
      Weight: 0, Metric Type: IGP
      15101 [Adjacency-SID, 13.1.1.1 - 13.1.1.2]
      15302
  Preference 600:
    Explicit: segment-list PE1-PE3-PE2 (inactive)
    Inactive Reason:
      Weight: 0, Metric Type: IGP
  Preference 500:
    Explicit: segment-list PE1-PE2 (inactive)
    Inactive Reason:
      Weight: 0, Metric Type: IGP
  Preference 200:
    Explicit: segment-list PE1-PE3-PE2-t0016 (inactive)
    Inactive Reason: unresolved first label
    Weight: 0, Metric Type: IGP
  Preference 100:
    Explicit: segment-list PE1-PE2_BE121 (inactive)
    Inactive Reason: unresolved first label
    Weight: 0, Metric Type: IGP
Attributes:
```

Associated Commands

```

Binding SID: 27525
  Allocation mode: dynamic
  State: Programmed
  Policy selected: yes
Forward Class: 0

```

```
PE1#show segment-routing traffic-eng forwarding policy name 100
```

```

Thu Feb 1 23:19:28.951 EST
Policy          Segment      Outgoing   Outgoing   Next Hop      Bytes
Name           List         Label      Interface  Next Hop      Switched
-----
100            PE1-PE2     Pop        Te0/0/0/2  12.1.9.2      0
                                   Pop          BE121      121.1.0.2    0

```

```
PE1#show segment-routing traffic-eng forwarding policy name 1013 detail
```

```

Thu Feb 1 21:22:46.069 EST
Policy          Segment      Outgoing   Outgoing   Next Hop      Bytes
Name           List         Label      Interface  Next Hop      Switched
-----
1013           PE1-PE3-PE2_link
                                   15302      Te0/0/0/1  13.1.1.2      0
Label Stack (Top -> Bottom): { 15302 }
Path-id: 1, Weight: 0
Packets Switched: 0
Local label: 24005
Packets/Bytes Switched: 0/0
(!): FRR pure backup

```

```
PE1#show mpls forwarding tunnels sr-policy name 1013
```

```

Thu Feb 1 21:23:22.743 EST
Tunnel          Outgoing   Outgoing   Next Hop      Bytes
Name           Label      Interface  Next Hop      Switched
-----
1013           (SR) 15302  Te0/0/0/1  13.1.1.2      0

```

Associated Commands

- [adjacency-sid](#)
- [index](#)
- [prefix-sid](#)
- [router isis](#)
- [segment-routing](#)

The applicable segment routing commands are described in the *Segment Routing Command Reference for Cisco NCS 5500, NCS 540 Series Routers, and NCS 560 Series Routers*.

Related Topics

- [Overview of Segment Routing , on page 203](#)
- [How Segment Routing Works , on page 204](#)
- [Segment Routing Global Block , on page 205](#)

EVPN VPWS On-Demand Next Hop with SR-TE

The EVPN VPWS On-Demand Next Hop with SR-TE feature enables you to fetch the best path to send traffic from the source to destination in a point-to-point service using IOS XR Traffic Controller (XTC). On-Demand Next Hop (ODN) with SR-TE is supported on EVPN Virtual Private Wire Service (VPWS) and Flexible Cross Connect (FXC) VLAN-unaware service.

When redistributing routing information across domains, provisioning of multi-domain services (Layer2 VPN and Layer 3 VPN) poses complexity and scalability issues. ODN with SR-TE feature delegates computation of an end-to-end Label Switched Path (LSP) to a path computation element (PCE). This PCE includes constraints and policies without any redistribution. It then installs the reapplied multi-domain LSP for the duration of the service into the local forwarding information base(FIB).

ODN uses BGP dynamic SR-TE capabilities and adds the path to the PCE. The PCE has the ability to find and download the end-to-end path based on the requirements. ODN triggers an SR-TE auto-tunnel based on the defined BGP policy. The PCE learns real-time topologies through BGP and/or IGP.

IOS XR Traffic Controller (XTC)

The path computation element (PCE) describes a set of procedures by which a path computation client (PCC) reports and delegates control of head-end tunnels sourced from the PCC to a PCE peer. The PCE peer requests the PCC to update and modify parameters of LSPs it controls. It also enables a PCC to allow the PCE to initiate computations and to perform network-wide orchestration.

Restrictions

- Maximum number of auto-provisioned TE policies is 1000.
- EVPN VPWS SR policy is not supported on EVPN VPWS dual homing.

EVPN validates if the route is for a single home next hop, otherwise it issues an error message about a dangling SR-TE policy, and continue to setup EVPN-VPWS without it. EVPN relies on ESI value being zero to determine if this is a single home or not. If the AC is a Bundle-Ether interface running LACP then you need to manually configure the ESI value to zero to overwrite the auto-sense ESI as EVPN VPWS multihoming is not supported.

To disable EVPN dual homing, configure bundle-Ether AC with ESI value set to zero.

```
evpn
interface Bundle-Ether12
  ethernet-segment
  identifier type 0 00.00.00.00.00.00.00.00
  /* Or globally */
evpn
ethernet-segment type 1 auto-generation-disable
```

Configure EVPN VPWS On Demand Next Hop with SR-TE

Perform the following steps to configure EVPN VPWS On Demand Next Hop with SR-TE. The following figure is used as a reference to explain the configuration steps:


- Configure Prefix-SID in ISIS

- Configure SR-TE
- Configure PCE and PCC
- Configure SR color
- Configure EVPN route policy
- Configure BGP
- Configure EVPN VPWS
- Configure Flexible Cross-connect Service (FXC) VLAN-unaware

Topology

Consider a topology where EVPN VPWS is configured on PE1 and PE2. Traffic is sent from PE1 to PE2 using SR-TE in the core. The PCE, which is configured on the P router, calculates the best path from PE1 to PE2. Path computation client (PCC) is configured on PE1 and PE2.

Figure 26: EVPN VPWS On Demand Next Hop with SR-TE



Configuration Example

Configure Prefix-SID in ISIS

Configure Prefix-SID in ISIS and topology-independent loop-free alternate path (TI-LFA) in the core such that each router uses a unique segment identifier associated with the prefix.

```
/* Configure Prefix-SID in ISIS and TI-LFA on PE1 */

Router# configure
Route(config)# router isis ring
Route(config-isis)# is-type level-2-only
Route(config-isis)# net 49.0001.1921.6800.1001.00
Route(config-isis)# segment-routing global-block 30100 39100
Route(config-isis)# nsr
Route(config-isis)# distribute link-state
Route(config-isis)# nsf cisco
Route(config-isis)# address-family ipv4 unicast
Route(config-isis-af)# metric-style wide
Route(config-isis-af)# mpls traffic-eng level-1
Route(config-isis-af)# mpls traffic-eng router-id loopback0
Route(config-isis-af)# segment-routing mpls
Route(config-isis-af)# exit
!
Route(config-isis)# interface loopback0
Route(config-isis-if)# address-family ipv4 unicast
Route(config-isis-af)# prefix-sid index 30101
Route(config-isis-af)# exit
!
Route(config-isis)# interface HundredGigE0/0/0/0
```



```

Route(config-isis-if) # circuit-type level-1
Route(config-isis-if) # point-to-point
Route(config-isis-if) # hello-padding disable
Route(config-isis-if) # fast-reroute per-prefix
Route(config-isis-if-af) # fast-reroute per-prefix ti-lfa
Route(config-isis-if-af) # commit

/*Configure Prefix-SID in ISIS and TI-LFA on P router */

Router# configure
Route(config) # router isis ring
Route(config-isis) # net 49.0001.1921.6800.1002.00
Route(config-isis) # segment-routing global-block 30100 39100
Route(config-isis) # nsr
Route(config-isis) # distribute link-state
Route(config-isis) # nsf cisco
Route(config-isis) # address-family ipv4 unicast
Route(config-isis-af) # metric-style wide
Route(config-isis-af) # mpls traffic-eng level-1
Route(config-isis-af) # mpls traffic-eng router-id loopback0
Route(config-isis-af) # segment-routing mpls
Route(config-isis-af) # exit
!
Route(config-isis) # interface loopback0
Route(config-isis-if) # address-family ipv4 unicast
Route(config-isis-af) # prefix-sid index 30102
Route(config-isis-af) # exit
!
Route(config-isis) # interface HundredGigE0/0/0/0
Route(config-isis-if) # circuit-type level-1
Route(config-isis-if) # point-to-point
Route(config-isis-if) # hello-padding disable
Route(config-isis-if) # fast-reroute per-prefix
Route(config-isis-if-af) # fast-reroute per-prefix ti-lfa
Route(config-isis-if-af) # exit
!
Route(config-isis) # interface HundredGigE0/0/0/1
Route(config-isis-if) # circuit-type level-1
Route(config-isis-if) # point-to-point
Route(config-isis-if) # hello-padding disable
Route(config-isis-if) # fast-reroute per-prefix
Route(config-isis-if-af) # fast-reroute per-prefix ti-lfa
Route(config-isis-if-af) # commit

/* Configure Prefix-SID in ISIS and TI-LFA on PE2 */

Router# configure
Route(config) # router isis ring
Route(config-isis) # net 49.0001.1921.6800.1003.00
Route(config-isis) # segment-routing global-block 30100 39100
Route(config-isis) # nsr
Route(config-isis) # distribute link-state
Route(config-isis) # nsf cisco
Route(config-isis) # address-family ipv4 unicast
Route(config-isis-af) # metric-style wide
Route(config-isis-af) # mpls traffic-eng level-1
Route(config-isis-af) # mpls traffic-eng router-id loopback0
Route(config-isis-af) # segment-routing mpls
Route(config-isis-af) # exit
!
Route(config-isis) # interface loopback0
Route(config-isis-if) # address-family ipv4 unicast

```

```

Route(config-isis-af) # prefix-sid index 30103
Route(config-isis-af) # exit
!
Route(config-isis) # interface HundredGigE0/0/0/1
Route(config-isis-if) # circuit-type level-1
Route(config-isis-if) # point-to-point
Route(config-isis-if) # hello-padding disable
Route(config-isis-if) # fast-reroute per-prefix
Route(config-isis-if-af) # fast-reroute per-prefix ti-lfa
Route(config-isis-if-af) # commit

```

Configure SR-TE

Configure SR-TE for P and PE routers.

```

/Configure SR-TE on PE1 */

Router# configure
Router(config) # segment-routing
Router(config-sr) # traffic-eng
Router(config-sr-te) # on-demand color 1
Router(config-sr-te-color) # dynamic mpls
Router(config-sr-te-color-dyn) # pcep
Router(config-sr-te-color-dyn-mpls-pce) # exit
!
Router(config-sr-te) # on-demand color 2
Router(config-sr-te-color) # dynamic mpls
Router(config-sr-te-color-dyn) # pcep
Router(config-sr-te-color-dyn-mpls-pce) # exit
!
Router(config-sr-te) # on-demand color 3
Router(config-sr-te-color) # dynamic mpls
Router(config-sr-te-color-dyn) # pcep
Router(config-sr-te-color-dyn-mpls-pce) # commit

/*Configure SR-TE on P router */
Router# configure
Router(config) # segment-routing
Router(config-sr) # traffic-eng
Router(config-sr-te) # commit

/Configure SR-TE on PE2 */

Router# configure
Router(config) # segment-routing
Router(config-sr) # traffic-eng
Router(config-sr-te) # on-demand color 11
Router(config-sr-te-color) # dynamic mpls
Router(config-sr-te-color-dyn) # pcep
Router(config-sr-te-color-dyn-mpls-pce) # exit
!
Router(config-sr-te) # on-demand color 12
Router(config-sr-te-color) # dynamic mpls
Router(config-sr-te-color-dyn) # pcep
Router(config-sr-te-color-dyn-mpls-pce) # exit
!
Router(config-sr-te) # on-demand color 13
Router(config-sr-te-color) # dynamic mpls
Router(config-sr-te-color-dyn) # pcep
Router(config-sr-te-color-dyn-mpls-pce) # commit

```

Configure PCE and PCC

Configure PCE on P router, and PCC on PE1 and PE2. Optionally, you can configure multiple PCEs as well.

```

/* Configure PCC on PE1 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc
Router(config-sr-te-pcc)# source-address ipv4 205.1.0.1
Router(config-sr-te-pcc)# pce address ipv4 205.2.0.2
Router(config-sr-te-pcc)# commit

/* Configure PCE on P router */

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# exit
Router(config)# pce
Router(config-pce)# address ipv4 205.2.0.2
Router(config-pce)# commit

/* Configure PCC on PE2 */

Router# configure
Router(config)# segment-routing
Router(config-sr)# traffic-eng
Router(config-sr-te)# pcc
Router(config-sr-te-pcc)# source-address ipv4 205.3.0.3
Router(config-sr-te-pcc)# pce address ipv4 205.2.0.2
Router(config-sr-te-pcc)# commit

```

Configure SR Color

Configure SR colors on PE routers.

```

/* Define SR color on PE1 */

Router# configure
Router(config)# extcommunity-set opaque color1
Router(config-ext)# 1
Router(config-ext)# end-set
!
Router(config)# extcommunity-set opaque color2
Router(config-ext)# 2
Router(config-ext)# end-set
!
Router(config)# extcommunity-set opaque color3
Router(config-ext)# 3
Router(config-ext)# end-set
!
/* Define SR color on PE2 */

Router# configure
Router(config)# extcommunity-set opaque color11
Router(config-ext)# 11
Router(config-ext)# end-set

```

```

!
Router(config)# extcommunity-set opaque color12
Router(config-ext)# 12
Router(config-ext)# end-set
!
Router(config)# extcommunity-set opaque color13
Router(config-ext)# 13
Router(config-ext)# end-set
!

```

Configure EVPN Route Policy

Configure EVPN route policy on PE1 and PE2. This example shows how to define the route policy language and track the EVPN route. The "rd" refers to the address of the PE and acts as Ethernet virtual interconnect for the L2 service.

```

/* Configure EVPN route policy on PE1 */

Router# configure
Router(config)# route-policy evpn_odn_policy
Router(config-rpl)# if rd in (205.3.0.3:2) then
Router(config-rpl-if)# set extcommunity color color1
Router(config-rpl-if)# set next-hop 205.3.0.3
Router(config-rpl-if)# elseif rd in (205.3.0.3:3) then
Router(config-rpl-elseif)# set extcommunity color color2
Router(config-rpl-elseif)# set next-hop 205.3.0.3
Router(config-rpl-elseif)# elseif rd in (205.3.0.3:4) then
Router(config-rpl-elseif)# set extcommunity color color3
Router(config-rpl-elseif)# set next-hop 205.3.0.3
Router(config-rpl-elseif)# endif
Router(config-rpl)# pass
Router(config-rpl)# end-policy

/* Configure EVPN route policy on PE2 */

Router# configure
Router(config)# route-policy evpn_odn_policy
Router(config-rpl)# if rd in (205.1.0.1:2) then
Router(config-rpl-if)# set extcommunity color color11
Router(config-rpl-if)# set next-hop 205.1.0.1
Router(config-rpl-if)# elseif rd in (205.1.0.1:3) then
Router(config-rpl-elseif)# set extcommunity color color12
Router(config-rpl-elseif)# set next-hop 205.1.0.1
Router(config-rpl-elseif)# elseif rd in (205.1.0.1:4) then
Router(config-rpl-elseif)# set extcommunity color color13
Router(config-rpl-elseif)# set next-hop 205.1.0.1
Router(config-rpl-elseif)# endif
Router(config-rpl)# pass
Router(config-rpl)# end-policy

```

Configure BGP

Configure BGP on PE1 and PE2.

```

/* Configure BGP on PE1 */

Router# configure
Router(config)# router bgp 100
Router(config-bgp)# bgp router-id 205.1.0.1

```

```

Routerconfig-bgp) # bgp graceful-restart
Router(config-bgp) # address-family l2vpn evpn
Router(config-bgp-af) # exit
!
Router(config-bgp) # neighbor 205.3.0.3
Router(config-bgp-nbr) # remote-as 100
Router(config-bgp-nbr) # update-source loopback 0
Router(config-bgp-nbr) # address-family l2vpn evpn
Router(config-bgp-nbr-af) # route-policy evpn_odn_policy in
Router(config-rpl) # commit

/* Configure BGP on PE2 */

Router# configure
Router(config) # router bgp 100
Routerconfig-bgp) # bgp router-id 205.3.0.3
Routerconfig-bgp) # bgp graceful-restart
Router(config-bgp) # address-family l2vpn evpn
Router(config-bgp-af) # exit
!
Router(config-bgp) # neighbor 205.1.0.1
Router(config-bgp-nbr) # remote-as 100
Router(config-bgp-nbr) # update-source loopback 0
Router(config-bgp-nbr) # address-family l2vpn evpn
Router(config-bgp-nbr-af) # route-policy evpn_odn_policy in
Router(config-rpl) # commit

```

Configure EVPN VPWS

Configure EVPN VPWS on PE1 and PE2.

```

/* Configure EVPN VPWS on PE1 */

Router# configure
Router(config) # interface GigE0/0/0/2.2 l2transport
Router(config-subif) # encapsulation dot1q 1
Router# exit
!
Router(config) # l2vpn
Router(config-l2vpn) # xconnect group evpn_vpws
Router(config-l2vpn-xc) # p2p e1_10
Router(config-l2vpn-xc-p2p) # interface GigE0/0/0/2.2
Router(config-l2vpn-xc-p2p) # neighbor evpn evi 2 target 10 source 10
Router(config-l2vpn-xc-p2p) # commit

/* Configure EVPN VPWS on PE2 */

Router# configure
Router(config) # interface GigE0/0/0/2.4 l2transport
Router(config-subif) # encapsulation dot1q 1
Router# exit
!
Router(config) # l2vpn
Router(config-l2vpn) # xconnect group evpn_vpws
Router(config-l2vpn-xc) # p2p e3_30
Router(config-l2vpn-xc-p2p) # interface GigE0/0/0/2.4
Router(config-l2vpn-xc-p2p) # neighbor evpn evi 2 target 10 source 10
Router(config-l2vpn-xc-p2p) # commit

```

Configure Flexible Cross-connect Service (FXC) VLAN-unaware

```

/* Configure FXC on PE1 */

Router# configure
Router(config)# interface GigE0/0/0/2.3 l2transport
Router(config-subif)# encapsulation dot1q 3
Router# exit
!
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware evpn_vu
Router(config-l2vpn-fxs-vu)# interface GigE0/0/0/2.3
Router(config-l2vpn-fxs-vu)# neighbor evpn evi 3 target 20
Router(config-l2vpn-fxs-vu)#commit

/* Configure FXC on PE2 */

Router# configure
Router(config)# interface GigE0/0/0/2.3 l2transport
Router(config-subif)# encapsulation dot1q 3
Router# exit
!
Router(config)# l2vpn
Router(config-l2vpn)# flexible-xconnect-service vlan-unaware evpn_vu
Router(config-l2vpn-fxs-vu)# interface GigE0/0/0/2.3
Router(config-l2vpn-fxs-vu)# neighbor evpn evi 3 target 20
Router(config-l2vpn-fxs-vu)#commit

```

Running Configuration

```

/* Configure Prefix-SID in ISIS and TI-LFA */

PE1:

configure
router isis ring
net 49.0001.1921.6800.1001.00
segment-routing global-block 30100 39100
nsr
distribute link-state
nsf cisco
address-family ipv4 unicast
metric-style wide
mpls traffic-eng level-1
mpls traffic-eng router-id Loopback0
segment-routing mpls
!
interface Loopback0
address-family ipv4 unicast
prefix-sid index 30101
!
!
interface HundredGigE0/0/0/0
circuit-type level-1
point-to-point
hello-padding disable
address-family ipv4 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
!

```

```

!
P:
configure
router isis ring
net 49.0001.1921.6800.1002.00
segment-routing global-block 30100 39100
nsr
distribute link-state
nsf cisco
address-family ipv4 unicast
metric-style wide
mpls traffic-eng level-1
mpls traffic-eng router-id Loopback0
segment-routing mpls
!
interface Loopback0
address-family ipv4 unicast
prefix-sid index 30102
!
!
interface HundredGigE0/0/0/0
circuit-type level-1
point-to-point
hello-padding disable
address-family ipv4 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
!
!
interface HundredGigE0/0/0/1
circuit-type level-1
point-to-point
hello-padding disable
address-family ipv4 unicast
fast-reroute per-prefix
fast-reroute per-prefix ti-lfa
!

```

PE2:

```

configure
router isis ring
net 49.0001.1921.6800.1003.00
segment-routing global-block 30100 39100
nsr
distribute link-state
nsf cisco
address-family ipv4 unicast
metric-style wide
mpls traffic-eng level-1
mpls traffic-eng router-id Loopback0
segment-routing mpls
!
interface Loopback0
address-family ipv4 unicast
prefix-sid index 30103
!
!
interface HundredGigE0/0/0/1
circuit-type level-1
point-to-point
hello-padding disable

```

```

address-family ipv4 unicast
  fast-reroute per-prefix
  fast-reroute per-prefix ti-lfa
!
!

/* Configure SR-TE */

PE1:

configure
segment-routing
traffic-eng
  on-demand color 1
  dynamic mpls
  pce
!
!
  on-demand color 2
  dynamic mpls
  pce
!
!
  on-demand color 3
  dynamic mpls
  pce
!
P:

configure
segment-routing
traffic-eng
!

PE2:

configure
segment-routing
traffic-eng
  on-demand color 11
  dynamic mpls
  pce
!
!
  on-demand color 12
  dynamic mpls
  pce
!
!
  on-demand color 13
  dynamic mpls
  pce
!

/* Configure PCE and PCC */

PE1:

configure
segment-routing
traffic-eng

```



```
    pcc
      source-address ipv4 205.1.0.1
      pce address ipv4 205.2.0.2
      !
```

PE:

```
configure
  segment-routing
  traffic-eng
  pce
    address ipv4 205.2.0.2
    !
```

PE2:

```
configure
  segment-routing
  traffic-eng
  pcc
    source-address ipv4 205.3.0.3
    pce address ipv4 205.2.0.2
    !
```

```
/* Configure SR Color */
```

PE1:

```
configure
  extcommunity-set opaque color1
  1
end-set
!
  extcommunity-set opaque color2
  2
end-set
!
  extcommunity-set opaque color3
  3
end-set
!
```

PE2:

```
configure
  extcommunity-set opaque color11
  11
end-set
!
  extcommunity-set opaque color12
  12
end-set
!
  extcommunity-set opaque color13
  13
end-set
!
```

```
/* Configure EVPN route policy */
```

PE1:

```
configure
  route-policy evpn_odn_policy
  if rd in (205.3.0.3:2) then
```

```

        set extcommunity color color1
        set next-hop 205.3.0.3
    elseif rd in (205.3.0.3:3) then
        set extcommunity color color2
        set next-hop 205.3.0.3
    elseif rd in (205.3.0.3:4) then
        set extcommunity color color3
        set next-hop 205.3.0.3
    endif
pass
end-policy

```

PE2:

```

configure
route-policy evpn_odn_policy
if rd in (205.1.0.1:2) then
    set extcommunity color color11
    set next-hop 205.1.0.1
elseif rd in (205.1.0.1:3) then
    set extcommunity color color12
    set next-hop 205.1.0.1
elseif rd in (205.1.0.1:4) then
    set extcommunity color color13
    set next-hop 205.1.0.1
endif
pass
end-policy

```

```
/* Configure BGP */
```

PE1:

```

configure
router bgp 100
  bgp router-id 205.1.0.1
  bgp graceful-restart
  address-family l2vpn evpn
  !
  neighbor 205.3.0.3
    remote-as 100
    update-source Loopback0
    address-family l2vpn evpn
    route-policy evpn_odn_policy in
  !

```

PE2:

```

configure
router bgp 100
  bgp router-id 205.3.0.3
  bgp graceful-restart
  address-family l2vpn evpn
  !
  neighbor 205.1.0.1
    remote-as 100
    update-source Loopback0
    address-family l2vpn evpn
    route-policy evpn_odn_policy in
  !

```

```
/* Configure EVPN VIEWS */
```

PE1:

```

configure
 interface GigE0/0/0/2.2 l2transport
  encapsulation dot1q 1
!
l2vpn
 xconnect group evpn_vpws
  p2p e1_10
  interface GigE0/0/0/2.2
    neighbor evpn evi 2 target 10 source 10
!
!

PE2:

configure
 interface GigE0/0/0/2.4 l2transport
  encapsulation dot1q 1
!
l2vpn
 xconnect group evpn_vpws
  p2p e3_30
  interface GigE0/0/0/2.4
    neighbor evpn evi 2 target 10 source 10
!
!
!

/* Configure Flexible Cross-connect Service (FXC) */

PE1:

configure
 interface GigE0/0/0/2.3 l2transport
  encapsulation dot1q 3
!
l2vpn
 flexible-xconnect-service vlan-unaware evpn_vu
 interface GigE0/0/0/2.3
  neighbor evpn evi 3 target 20
!
!

PE2:

configure
 interface GigE0/0/0/2.3 l2transport
  encapsulation dot1q 3
!
l2vpn
 flexible-xconnect-service vlan-unaware evpn_vu
 interface GigE0/0/0/2.3
  neighbor evpn evi 3 target 20
!
!

```

Verify EVPN VPWS On Demand Next Hop with SR-TE Configuration

Verify if SR-TE policy is auto-provisioned for each L2 service configured on EVPN ODN.

```
PE1# show segment-routing traffic-eng policy
```

SR-TE policy database

Name: **bgp_AP_1 (Color: 1, End-point: 205.3.0.3)**

Status:

Admin: up Operational: up for 07:16:59 (since Oct 3 16:47:04.541)

Candidate-paths:

Preference 100:

Dynamic (pce 205.2.0.2) (active)

Weight: 0

30103 [Prefix-SID, 205.3.0.3]

Attributes:

Binding SID: 68007

Allocation mode: dynamic

State: Programmed

Policy selected: yes

Forward Class: 0

Distinguisher: 0

Auto-policy info:

Creator: BGP

IPv6 caps enable: no

PE1#show l2vpn xconnect group evpn_vpws xc-name evpn_vpws_1001 detail

Group evpn_vpws, XC evpn_vpws_1001, state is up; Interworking none

AC: Bundle-Ether12.1001, state is up

Type VLAN; Num Ranges: 1

Outer Tag: 1000

Rewrite Tags: []

VLAN ranges: [1, 1]

MTU 1500; XC ID 0xc0000018; interworking none

Statistics:

packets: received 642304, sent 642244

bytes: received 61661184, sent 61655424

drops: illegal VLAN 0, illegal length 0

EVPN: neighbor 20.20.20.20, PW ID: evi 1001, ac-id 10001, state is up (established)

XC ID 0xa0000007

Encapsulation MPLS

Source address 10.10.10.10

Encap type Ethernet, control word enabled

Sequencing not set

Preferred path Active : SR TE pref_sid_to_PE3, On-Demand, fallback enabled

Tunnel : Up

Load Balance Hashing: src-dst-mac

PE1#show bgp l2vpn evpn route-type 1

BGP router identifier 205.1.0.1, local AS number 100

BGP generic scan interval 60 secs

Non-stop routing is enabled

BGP table state: Active

Table ID: 0x0 RD version: 0

BGP main routing table version 36

BGP NSR Initial initsync version 25 (Reached)

BGP NSR/ISSU Sync-Group versions 36/0

BGP scan interval 60 secs

Status codes: s suppressed, d damped, h history, * valid, > best

i - internal, r RIB-failure, S stale, N Nexthop-discard

Origin codes: i - IGP, e - EGP, ? - incomplete

Network Next Hop Metric LocPrf Weight Path

Route Distinguisher: 205.1.0.1:2 (default for vrf VPWS:2)

*>i[1][0000.0000.0000.0000.0000][1]/120

205.3.0.3 T:bgp_AP_1

100 0 i

```
PE1# show evpn evi ead detail

EVI Ethernet Segment Id EtherTag Nexthop Label SRTE IFH
-----
-----
2 0000.0000.0000.0000.0000 1 205.3.0.3 24000 0x5a0
Source: Remote, MPLS
```

Associated Commands

- adjacency-sid
- index
- prefix-sid
- [router isis](#)
- segment-routing

The applicable segment routing commands are described in the *Segment Routing Command Reference for Cisco NCS 5500 Series Routers, Cisco NCS 540 Series Routers, and Cisco NCS 560 Series Routers*.

Related Topics

- [Overview of Segment Routing](#) , on page 203
- [How Segment Routing Works](#) , on page 204
- [Segment Routing Global Block](#) , on page 205

Overview of Segment Routing

Segment Routing (SR) is a flexible, scalable way of doing source routing. The source chooses a path and encodes it in the packet header as an ordered list of segments. Segments are identifier for any type of instruction. Each segment is identified by the segment ID (SID) consisting of a flat unsigned 32-bit integer. Segment instruction can be:

- Go to node N using the shortest path
- Go to node N over the shortest path to node M and then follow links Layer 1, Layer 2, and Layer 3
- Apply service S

With segment routing, the network no longer needs to maintain a per-application and per-flow state. Instead, it obeys the forwarding instructions provided in the packet.

Segment Routing relies on a small number of extensions to Cisco Intermediate System-to-Intermediate System (IS-IS) and Open Shortest Path First (OSPF) protocols. It can operate with an MPLS (Multiprotocol Label Switching) or an IPv6 data plane, and it integrates with the rich multi service capabilities of MPLS, including Layer 3 VPN (L3VPN), Virtual Private Wire Service (VPWS), and Ethernet VPN (EVPN).

Segment routing can be directly applied to the Multiprotocol Label Switching (MPLS) architecture with no change in the forwarding plane. Segment routing utilizes the network bandwidth more effectively than traditional MPLS networks and offers lower latency. A segment is encoded as an MPLS label. An ordered

list of segments is encoded as a stack of labels. The segment to process is on the top of the stack. The related label is popped from the stack, after the completion of a segment.

Segment Routing provides automatic traffic protection without any topological restrictions. The network protects traffic against link and node failures without requiring additional signaling in the network. Existing IP fast re-route (FRR) technology, in combination with the explicit routing capabilities in Segment Routing guarantees full protection coverage with optimum backup paths. Traffic protection does not impose any additional signaling requirements.

How Segment Routing Works

A router in a Segment Routing network is capable of selecting any path to forward traffic, whether it is explicit or Interior Gateway Protocol (IGP) shortest path. Segments represent subpaths that a router can combine to form a complete route to a network destination. Each segment has an identifier (Segment Identifier) that is distributed throughout the network using new IGP extensions. The extensions are equally applicable to IPv4 and IPv6 control planes. Unlike the case for traditional MPLS networks, routers in a Segment Router network do not require Label Distribution Protocol (LDP) and Resource Reservation Protocol - Traffic Engineering (RSVP-TE) to allocate or signal their segment identifiers and program their forwarding information.

There are two ways to configure segment routing:

- SR-TE policy under "segment-routing traffic-eng" sub-mode
- TE tunnel with SR option under "mpls traffic-eng" sub-mode



Note

However, you can configure the above mentioned L2VPN and EVPN services using only "segment-routing traffic-eng" sub-mode.

Each router (node) and each link (adjacency) has an associated segment identifier (SID). Node segment identifiers are globally unique and represent the shortest path to a router as determined by the IGP. The network administrator allocates a node ID to each router from a reserved block. On the other hand, an adjacency segment ID is locally significant and represents a specific adjacency, such as egress interface, to a neighboring router. Routers automatically generate adjacency identifiers outside of the reserved block of node IDs. In an MPLS network, a segment identifier is encoded as an MPLS label stack entry. Segment IDs direct the data along a specified path. There are two kinds of segment IDs:

- Prefix SID: A segment ID that contains an IP address prefix calculated by an IGP in the service provider core network. Prefix SIDs are globally unique. A prefix segment represents the shortest path (as computed by IGP) to reach a specific prefix; a node segment is a special prefix segment that is bound to the loopback address of a node. It is advertised as an index into the node specific SR Global Block or SRGB.
- Adjacency SID: A segment ID that contains an advertising router's adjacency to a neighbor. An adjacency SID is a link between two routers. Since the adjacency SID is relative to a specific router, it is locally unique.

A node segment can be a multi-hop path while an adjacency segment is a one-hop path.

Segment Routing Global Block

Segment Routing Global Block (SRGB) is the range of labels reserved for segment routing. SRGB is local property of an segment routing node. In MPLS, architecture, SRGB is the set of local labels reserved for global segments. In segment routing, each node can be configured with a different SRGB value and hence the absolute SID value associated to an IGP Prefix Segment can change from node to node.

The SRGB default value is 16000 to 23999. The SRGB can be configured as follows:

```
Router(config)# router isis 1
Router(config-isis)#segment-routing global-block 45000 55000
```




CHAPTER 11

MSTP BPDU Guard

The MSTP BPDU Guard feature protects against misconfiguration of edge ports.



Note In order to enable the MSTP BPDU Guard feature for an interface, the command **portfast bpduguard** must be configured on it.

Port Fast

The Port Fast feature manage the ports at the edge of the switched Ethernet network. For devices that only have one link to the switched network (typically host devices), there is no need to run MSTP, as there is only one available path. Furthermore, it is undesirable to trigger topology changes (and resultant MAC flushes) when the single link fails or is restored, as there is no alternative path.

By default, MSTP monitors ports where no BPDUs are received, and after a timeout, places them into edge mode whereby they do not participate in MSTP. When **portfast** is explicitly configured on an interface, MSTP considers that interface to be an edge port and removes it from consideration when calculating the spanning tree. And hence the convergence time for the whole network is improved when **portfast** is configured.



Note MSTP BPDU Guard feature is supported by configuring interfaces in port fast mode. BPDU guard feature will error-disable the port on receiving BPDU packets.

- [Configuring MSTP BPDU Guard, on page 207](#)

Configuring MSTP BPDU Guard

This section describes how you can configure MSTP BPDU Guard.

```
Router# configure
Router(config)# l2vpn bridge group bg1
Router(config-l2vpn-bg)# bridge-domain bd1
Router(config-l2vpn-bg-bd)# int TenGigE 0/0/0/7
Router(config-l2vpn-bg-bd-ac)# root
Router(config)# spanning-tree mst m0
Router(config-mstp)# interface tenGigE 0/0/0/7
Router(config-mstp-if)# portfast bpduguard
```

```

Router(config-mstp-if)# root
Router(config)# int tenGigE 0/0/0/7 l2transport
Router(config-if-l2)# commit

```

Running Configuration with MSTP BPDU Guard

```

!
Configure
l2vpn
bridge group bg1
  bridge-domain bd1
  interface TenGigE0/0/0/7
  !
spanning-tree mst m0
  interface TenGigE0/0/0/7
  portfast bpduguard
!
interface TenGigE0/0/0/7
  l2transport
!

```

Verification for MSTP BPDU Guard

Verify that you have configured MSTP BPDU Guard.

```

/* Verify the MSTP BPDU Guard configuration */
Router# show interfaces tenGigE 0/0/0/7
Wed Nov  9 09:23:56.268 UTC
TenGigE0/0/0/7 is error disabled, line protocol is administratively down
  Interface state transitions: 2
  Hardware is TenGigE, address is 7cad.7425.c8c8 (bia 7cad.7425.c8c8)
  Layer 2 Transport Mode
  MTU 1514 bytes, BW 10000000 Kbit (Max: 10000000 Kbit)
    reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation ARPA,
  Full-duplex, 10000Mb/s, link type is force-up
  output flow control is off, input flow control is off
  Carrier delay (up) is 10 msec
  loopback not set,
  Last link flapped 00:00:49
  Last input 00:00:40, output 00:00:40
  Last clearing of "show interface" counters never
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    38752 packets input, 4611429 bytes, 0 total input drops
    1 drops for unrecognized upper-level protocol
  Received 1 broadcast packets, 38751 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort

```