



OcNOS®

**Open Compute Network Operating System
for Service Providers**

Key Features

Version 7.0.0

February 2026

©2026 IP Infusion Inc. All Rights Reserved.

This documentation is subject to change without notice. The software described in this document and this documentation are furnished under a license agreement or nondisclosure agreement. The software and documentation may be used or copied only in accordance with the terms of the applicable agreement. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's internal use without the written permission of IP Infusion Inc.

IP Infusion Inc.

3979 Freedom Circle, Suite 900

Santa Clara, CA 95054

+1 408-400-1900

<http://www.ipinfusion.com/>

For support, questions, or comments via E-mail, contact:

support@ipinfusion.com

Trademarks:

IP Infusion and OcNOS are trademarks or registered trademarks of IP Infusion. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Use of certain software included in this equipment is subject to the IP Infusion, Inc. End User License Agreement at <http://www.ipinfusion.com/license>. By using the equipment, you accept the terms of the End User License Agreement.

| CONTENTS

Contents	3
Preface	9
About this Guide	9
Audience	9
Conventions	9
IP Infusion Product Release Version	9
Related Documentation	10
Feature Availability	10
Migration Guide	10
IP Maestro Support	10
Technical Support	10
Technical Sales	10
Technical Documentation	10
Documentation Disclaimer	11
Comments	11
About This Guide	12
Overview	12
Target Users	12
Key Capabilities	12
Quick Links	12
Routing and Security Enhancements	13
BGP RPKI-Based Route Validation	15
Overview	15
Feature Characteristics	15
Benefits	15
Configuration	16
Prerequisites	16
Topology	16
Validation	18
Implementation Examples	20
Commands	20
bgp rpki server	21
bgp origin-as validation-enable	23
bgp origin-as bestpath	24
match rpki	25
show bgp rpki table ipv4	26
show bgp rpki table ipv6	28
show bgp origin-as validity ipv4	30
show bgp origin-as validity ipv6	31
show bgp rpki server	33

Glossary	34
ECMP Support for L3EVPN	35
Overview	35
Feature Characteristics	35
Benefits	36
Limitations	36
Supported Hardware	36
Prerequisites	36
Configuration	37
Topology	37
PE1	37
P1/RR1	38
P2	39
PE3	40
PE4	40
Configuration Snapshot	41
P1/RR1	55
P2/RR2	61
PE3	65
PE4	79
Validation	91
ECMP for L3EVPN Commands	98
multi-path	98
Command Syntax	98
Parameters	98
Default	98
Command Mode	98
Examples	99
EVPN AF Route-Maps and Route Filtering	100
Overview	100
Feature Characteristics	100
Benefits	100
Prerequisites	100
Limitation	101
Configuration	101
Validation	103
Implementation Examples	110
Commands	110
mac-list	111
match mac address list	112
match evpn-route-type	113
show mac-list	114
Glossary	114
BGP Labeled Unicast Next Hop Self in Route-Map	115
Overview	115
Feature Characteristics	115
Limitations	115

Prerequisites	115
Configuration	115
Configuration Snapshot	117
Commands	175
set ip next-hop self	175
Layer 2 Service Enhancements	176
EVPN MPLS E-Tree Scenario 2	177
Overview	177
Inter-PE E-tree: New Route Advertisements	177
Feature Characteristics	177
Benefits	179
E-Tree Scenario 2 Prerequisites	179
Root PEs and Leaf PEs	179
MPLS Core Fabric Formation (P1 or P2 Routers)	181
Access Switch VLAN Trunk Configuration	182
E-Tree Scenario 2 Configuration	182
Topology	182
Enable E-Tree Scenario-2 Functionality	183
VPN ID and Leaf Role Mapping	184
E-Tree Scenario 2 Validation	184
E-Tree Scenario 2 Network Topology Running Configurations	188
E-Tree Scenario 2 Implementation Examples	207
E-Tree Scenario 2 Troubleshooting	207
E-Tree Scenario 2 Revised Commands	208
evpn etree	209
map vpn-id	210
E-Tree Scenario 2 Glossary	210
Network Management and Automation	213
Container Runtime and Life cycle Management Using K3S	214
Overview	214
Feature Characteristics	214
Benefits	214
Prerequisites	215
Limitations	215
Configuration	215
Implementation Examples	217
K3s Containers Commands	220
feature k3s	220
feature docker	220
show feature k3s	221
show feature docker	222
Glossary	222
NetConf Access Control Model User Guide	224
Overview	224
Feature Characteristics	224
Configuration	225

Initial NACM Configuration	225
Prerequisites	226
Common NACM Rule Fields	226
Creating NetConf RPC for NACM	229
RPC Configurations for NACM	229
Glossary	238
sFlow - Sample Packet Monitoring for Multiple Interfaces	239
Overview	239
Features Characteristics	240
Benefits	240
Prerequisites	240
Configuration	241
Topology	241
Validation	243
Configuring sFlow with User Defined VRFs	243
Implementation Examples	246
Commands	252
no sflow collector-id	253
sflow collector	254
Troubleshooting	255
Glossary	255
VxLAN OAM for Overlay Networks	256
CLI-Script and CLI-Shell	257
Overview	257
Limitations	257
Configuration	257
Validation	258
CLI-Script and CLI-Shell Commands	261
System Limits and Counters	274
Overview	274
Feature Characteristics	274
Benefits	274
System Limits and Counters Limitation	275
Data Characteristics	275
Security and Access	275
Dependencies	275
Platform-specific	275
Encoding	275
System Limits and Counters Configuration	275
Topology	276
Use Case: Verify Total Number of IPv4 Routes Installed	276
System Limits and Counters Implementation Example	282
Scenario 1: Resource Audit During Large-Scale Migration	282
Scenario 2: Monitoring via Network Management Systems	282
Scenario 3: Capacity Pre-Check Before Service Deployment	282
System Limits and Counters Commands	282
show hardware-routing limits	282

show hardware-l2vpn-instances limits	283
System Limits and Counters Revised Commands	283
show hardware-routing limits	283
show hardware-l2vpn-instances limits	285
show access-lists	286
show ip vrf	287
show ip route	288
show ipv6 route	292
show interface	294
System Limit Counters Troubleshooting	297
Show Output Displays Blank or Partial Results	297
gNMI Returns “Unsupported Encoding” or Missing Fields	297
API Retrieval Fails for Specific Resource Paths	297
System Limit Counters Glossary	298
Layer 2 or Layer 3 Overlay Networking	300
Flow-Aware Transport for EVPN over MPLS	301
Overview	301
Feature Characteristics	301
Benefits	301
FAT Prerequisites	301
PE Nodes	302
FAT Configuration	305
Topology	305
PE Nodes Configuration	307
FAT Validation	308
FAT Implementation Example	311
FAT Revised Command	312
evpn mpls id	313
FAT Glossary	314
Recursive Next-hop Resolution with MPLS Next-hop	316
Overview	316
Feature Characteristics	316
Benefits	316
Limitation	316
Configuration	316
Configuration with MPLS LDP	316
Configuration with Segment Routing MPLS (SR-MPLS)	321
On PE1	326
On PE2	327
From PE1 to CE2	328
From PE2 to CE1	328
Implementation Example	329
Use Case: Simplified MPLS Core for BGP-Free Operation	329
Revised Commands	329
show ip route mpls	329
show ip route database	329

show rib txlist	329
Troubleshooting	330
show rib txlist	331
Glossary	332

PREFACE

About this Guide

This guide describes how to configure Key Features in OcNOS.


Audience

This guide is intended for network administrators and other engineering professionals who configure OcNOS.

Conventions

The [Table 1](#) table shows the conventions used in this guide.

Table 1. Conventions

Convention	Description
Italics	Emphasized terms; titles of books
 Note:	Special instructions, suggestions, or warnings
<code>monospaced type</code>	Code elements such as commands, parameters, files, and directories

IP Infusion Product Release Version

Each integer in release numbers indicates Major, Minor, and Maintenance release versions. Build numbers that follow the release numbers are for internal tracking and verification of the software build process and are visible to customers as part of the software version number.



Product Name: IP Infusion Product Family

Major Version: New customer-facing functionality that represents a significant change to the code base; including a significant marketing change or direction in the product.

Minor Version: Enhancements or extensions to existing features, changes to address external needs, or internal improvements to satisfy new sales regions or marketing initiatives.

Maintenance Version: A collection of product bugs or issues usually scheduled every 30 or 60 days, based on the number of issues.

Related Documentation

For information about installing OcNOS, see the *Installation Guide* for your platform.

Feature Availability

Each OcNOS SKU contains a set of supported features. For a list of available features based on the SKU that you purchased, refer to the [Feature Matrix](#).

Migration Guide

Check the *Migration Guide* for necessary configuration changes before migrating from one version of OcNOS to another.

IP Maestro Support

Monitor devices running OcNOS Release 6.3.4-70 and above using IP Maestro software.

Technical Support

IP Infusion maintains an online technical support site that provides a variety of technical support programs for licensed OcNOS customers at the [Technical Assistance Center](#).

Customers and partners enjoy full access to the support website. The site allows customers and partners to open technical support calls, update open calls with new information, and review the status of open or closed calls. The password-protected site includes technical documentation, Release Notes, and descriptions of service offerings.

Technical Sales

Contact the IP Infusion sales representative for more information about the OcNOS solution.

Technical Documentation

For core commands and configuration procedures, visit: [Product Documentation](#).

For training videos, visit: [OcNOS Free Training Videos](#).

For a list of supported platforms and SKUs of OcNOS features, refer to the [OcNOS Feature Matrix](#).

Documentation Disclaimer

The global documentation site is evolving to provide an enhanced website user experience for select topics included in this release. Some guides are now available outside the existing documentation library and can be accessed directly from custom documentation landing pages. These guides offer robust in-built search functionality.

For the latest documentation, visit the product-specific documentation landing page and select the relevant guide.

Comments

If you have comments, or need to report a problem with the content, contact techpubs@ipinfusion.com.

ABOUT THIS GUIDE

Overview

The OcNOS Key Features Guide provides a unified view of the most important enhancements and features introduced in each OcNOS release. It serves as a single reference for customers to quickly identify and understand the main capabilities delivered across different functional modules.

Unlike module-specific documentation, which includes detailed descriptions and high-level sections such as an overview, benefits, configuration procedures, implementation examples, and command details for deployment, this guide highlights only the key features introduced in each release. Each feature listed here is also documented in detail in the respective module guide.

Target Users

This guide is intended for network architects, operations teams, IT support engineers, product managers, and technical evaluators who need a consolidated view of new and enhanced capabilities introduced in each OcNOS release.

Key Capabilities

This guide enables users to:

- Track major feature additions and enhancements across releases in a single reference
- Compare capabilities introduced in different versions of OcNOS
- Quickly assess which new features align with deployment requirements
- Navigate to the corresponding module guides for detailed configuration and operational procedures

By presenting a structured, release-by-release summary of key capabilities, the Key Features guide helps users stay aligned with OcNOS developments, evaluate new capabilities at a glance, and efficiently locate the detailed documentation needed for planning, deployment, and support.

Quick Links

- [Routing and Security Enhancements \(page 13\)](#)
- [Layer 2 Service Enhancements \(page 176\)](#)
- [Network Management and Automation \(page 213\)](#)
- [Layer 2 or Layer 3 Overlay Networking \(page 300\)](#)

| ROUTING AND SECURITY ENHANCEMENTS

OcNOS continues to strengthen its Layer 3 routing and security capabilities, delivering improved route control, visibility, and policy enforcement across BGP, MPLS, and VRF environments. These updates enhance interoperability, stability, and operational flexibility in large-scale routed networks.

BGP RPKI-Based Route Validation	15
Overview	15
Feature Characteristics	15
Benefits	15
Configuration	16
Implementation Examples	20
Commands	20
Glossary	34
ECMP Support for L3EVPN	35
Overview	35
Prerequisites	36
Configuration	37
Topology	37
Configuration Snapshot	41
P1/RR1	55
P2/RR2	61
PE3	65
PE4	79
Validation	91
ECMP for L3EVPN Commands	98
multi-path	98
EVPN AF Route-Maps and Route Filtering	100
Overview	100
Feature Characteristics	100
Benefits	100
Prerequisites	100
Configuration	101
Implementation Examples	110
Commands	110
Glossary	114
BGP Labeled Unicast Next Hop Self in Route-Map	115
Overview	115
Feature Characteristics	115
Limitations	115

Prerequisites	115
Configuration	115
Commands	175

BGP RPKI-Based Route Validation

Overview

Resource Public Key Infrastructure (RPKI) is a security framework designed to mitigate the risk of BGP prefix hijacking by cryptographically verifying that an Autonomous System (AS) is authorized to announce a given IP prefix.

In OcNOS, RPKI-based BGP Origin Validation allows the router to download Route Origin Authorizations (ROAs) from an RPKI server via the RTR protocol. The downloaded ROAs are then used to validate incoming BGP routes, ensuring that only legitimate prefixes are considered during best path selection.

This feature improves routing security by reducing the acceptance and propagation of invalid routes.

Feature Characteristics

- **ROA Retrieval:** Supports downloading ROAs from multiple (up to 10) RPKI servers over TCP or SSH transport.
- **Per-AF and Per-VRF Support:** Validation can be enabled on a per-address-family (IPv4/IPv6 unicast) and per-VRF basis.
- **Validation States:** Each route is tagged with one of the three validation state:
 - **Valid (V):** Prefix-AS match found in ROA.
 - **Invalid (I):** Prefix-AS mismatch or not authorized.
 - **Not-Found (N):** No corresponding ROA.
- **Flexible Policy Control:** Route-map support for matching on RPKI state (valid, invalid, not-found) to set attributes such as local preference.
- **Best Path Selection Control:**
 - Option to consider only valid/not-found routes for path selection.
 - Configurable to allow invalid routes in best path preference.
- **Dynamic Updates:** ROA updates are applied in real time from RPKI servers.

Benefits

- **Enhanced Security:** Prevents acceptance of hijacked or misconfigured routes.
- **Operational Flexibility:** Operators can tune route selection with route-maps or allow invalid routes for troubleshooting.
- **Standards Compliance:** Implements BGP Origin Validation as per RPKI-based validation standards.
- **Granular Control:** Policies can be applied per Address Family (AF) or Virtual Routing and Forwarding (VRF), giving operators flexibility in deploying validation gradually.
- **Improved Resilience:** Reduces propagation of invalid prefixes across the Internet routing system.

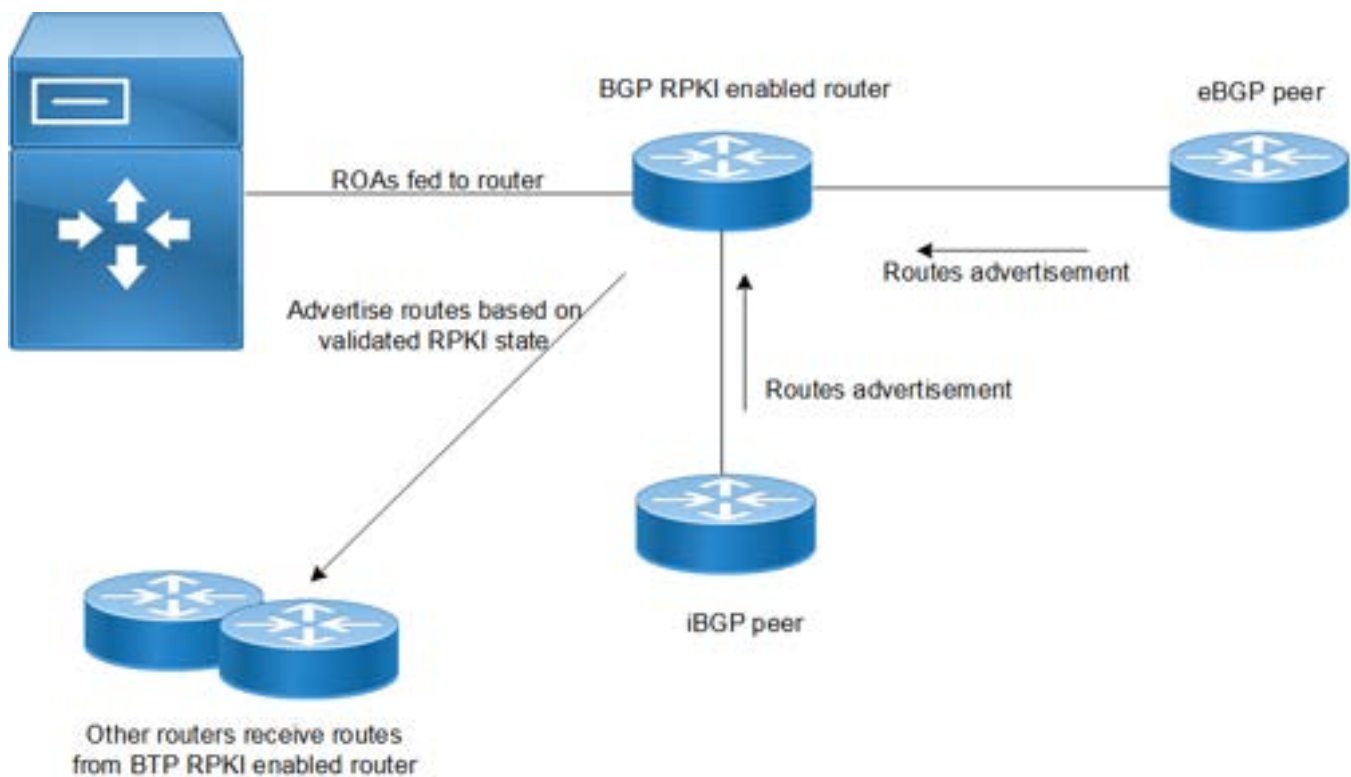
Configuration

This section describes the configuration procedure for enabling RPKI-based BGP origin validation.

Prerequisites

- BGP is already configured (router bgp (ASN)) on the device.
- You have reachable RPKI RTR server IP(s) and credentials (if using SSH).
- Make sure the required transport ports (TCP or SSH) are reachable between the OcNOS and the RPKI server(s).
- Decide per address-family (IPv4/IPv6) and per VRF where you want validation turned on.

Topology



In this topology, the RPKI Validator stores and distributes Route Origin Authorizations (ROAs), which contain information about the prefixes and its authorized originating AS numbers. The validator communicates these validated ROAs to the BGP RPKI-enabled router.

The BGP RPKI-enabled router establishes a session with the RPKI Validator using either TCP or SSH, depending on the configuration. Upon receiving ROAs, the router validates the BGP route advertisements against the authorized prefix–origin pairs.

Based on the validation results, the router marks each route as Valid, Invalid, or NotFound, and applies routing policies accordingly.

Only routes that pass the validation check (Valid) are used for forwarding or are advertised to other peers.

- ROAs fed to router: The RPKI Validator sends validated prefix-origin data to the BGP RPKI-enabled router.

- Advertise routes based on validated RPKI state: The router advertises only validated routes to its BGP peers.
- iBGP and eBGP peers: Both internal (iBGP) and external (eBGP) peers receive route updates from the RPKI-enabled router.
- Other routers: Routers within the same AS or network domain receive routes that have already been validated, ensuring route authenticity and preventing prefix hijacking.

1. Configure one or more RPKI servers to establish RTR sessions and download Route Origin Authorizations (ROAs).

```
ocnos(config)# router bgp 100
ocnos(config-router)# bgp rpki server 10.30.0.85 ssh user test encrypt 0 password test refresh 1
retry 1 expire 600
ocnos(config-router)# bgp rpki server 192.168.1.233 tcp refresh 1 retry 1 expire 600
ocnos(config-router)# commit
```

2. Enable origin validation on the required address family (AF) or VRF. This allows BGP routes to be tagged with an RPKI validation state:

- V: Valid
- I: Invalid
- N: Not-found

```
ocnos(config)# router bgp 100
ocnos(config-router)# address-family ipv4 unicast
ocnos(config-router-af)# bgp origin-as validation-enable
ocnos(config-router-af)# commit
```

3. Configure BGP to consider RPKI validation state in the best-path selection process. Invalid routes are excluded, and preference is given in the following order: valid > not-found (unless modified by policy).

```
ocnos(config-router-af)# bgp origin-as bestpath use-validity
ocnos(config-router-af)# commit
```

4. Permit invalid routes to participate in best-path selection but assign them the lowest preference.

```
ocnos(config-router-af)# bgp origin-as bestpath allow-invalid
ocnos(config-router-af)# commit
```

5. Use route-maps to define policy actions, such as setting local preference, based on the RPKI validation state of a route.

Example of the route map:

```
route-map RPKI-1 permit 3
match rpki valid
set local-preference 100

route-map RPKI-1 permit 5
match rpki not-found
set local-preference 200

route-map RPKI-1 permit 10
match rpki invalid
set local-preference 300

ocnos(config-router)# address-family ipv4 unicast
ocnos(config-router-af)# neighbor 100.1.1.2 activate
ocnos(config-router-af)# neighbor 100.1.1.2 route-map RPKI-1 in
ocnos(config-router-af)# commit
```

6. Remove or rollback RPKI configuration:

- Disable Validation in an AF/VRF:

```
ocnos(config-router-af)# no bgp origin-as validation-enable
ocnos(config-router-af)# commit
```

- Remove an RPKI Server:

```
ocnos(config)# router bgp 100
ocnos(config-router)# no bgp rpki server 10.30.0.85
ocnos(config-router)# commit
```

Validation

Verify server session for the following:

State: Established and Synced: TRUE in show bgp rpki server detail.

```
OCNOS#show bgp rpki server detail
BGP RPKI Server Information
  Server Address: 155.155.1.1:
    Transport: TCP:3323
    RTR Version: 1
    State: established
    Synced: TRUE
    Uptime: 00:00:39
    ROAs (IPv4/IPv6): 9/3
    Configured Refresh-Interval: 15 seconds
    Configured Retry-Interval: 10 seconds
    Configured Expire-Interval: 600 seconds
    Actual Refresh-Interval: 5 seconds
    Actual Retry-Interval: 5 seconds
    Actual Expire-Interval: 600 seconds
    Rest of time to Refresh-Interval expiration: 2 seconds
    Rest of time to Expire-Interval expiration: 597 seconds
    ToBeDeleted: FALSE
```

Verify validation state on routes

```
ocnos# show bgp origin-as validity ipv4
BGP table version is 28, local router ID is 1.1.1.1

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,

                l - labeled, S Stale, x-EVPN

Origin codes: i - IGP, e - EGP, ? - incomplete

Description : Ext-Color - Extended community color

Origin-AS validation codes: V valid, I invalid, N not-found, D disabled
```

	Network	Next Hop	Metric	LocPrf	Weight	Path	Ext-Color
N*>	1.2.0.0/16	0.0.0.0	0	100	32768	?	-
I*>	1.2.11.0/24	0.0.0.0	0	100	32768	i	-
I*		100.1.1.7	0	100	32768	?	-
N*	1.2.21.30/32	0.0.0.0	0	100	32768	i	-
N*>	1.6.0.0/16	0.0.0.0	0	100	32768	{300,9583} ?	-

I*>i	1.6.14.0/24	100.1.1.7	0	100	0	?	-	
I*		100.1.1.7	20	100	32768	?	-	
V*>i	1.6.136.0/24	100.4.1.5	0	100	0	9583	?	-
V*		100.2.1.3	0	100	0	300	9583	?
N*								-
N*>i	2.0.0.1/32	0.0.0.0	0	100	32768	i	-	
N*>i	2.2.2.2/32	100.1.1.2	0	100	0	?	-	
I*	2.3.4.5/32	0.0.0.0	0	100	32768	i	-	
I*>	3.3.3.3/32	100.2.1.3	0	100	0	300	?	-
I*	3.4.5.6/32	0.0.0.0	0	100	32768	i	-	
N*>i	4.4.4.4/32	100.4.1.4	0	100	0	400	?	-
N*		100.2.1.3	0	100	0	300	400	?
I*>i	5.5.5.5/32	100.4.1.5	0	100	0	9583	?	-
I*		100.2.1.3	0	100	0	300	9583	?
N*>i	7.7.7.7/32	100.1.1.7	0	100	0	?	-	
N*		100.1.1.7	20	100	32768	?	-	
I*>	8.8.8.8/32	100.1.1.8	0	100	0	400	1300	?
N*	33.44.55.66/32	0.0.0.0	0	100	32768	i	-	
N*>	100.1.1.0/24	100.1.1.8	0	100	0	400	1300	?
N* i		100.1.1.2	0	100	0	?	-	
N* i		100.1.1.7	0	100	0	?	-	
N*		0.0.0.0	1	100	32768	?	-	
I*>	100.2.1.0/24	100.2.1.3	0	100	0	300	?	-
I*		100.1.1.8	0	100	0	400	1300	?
I* i		100.1.1.7	0	100	0	?	-	
N*>	100.3.1.0/24	100.2.1.3	0	100	0	300	?	-
I*>	100.4.1.0/24	100.2.1.3	0	100	0	300	400	?
I* i		100.1.1.2	0	100	0	?	-	
N*>	172.16.181.0/24	100.2.1.3	0	100	0	300	?	-
N*		100.1.1.8	0	100	0	400	1300	?
N* i		100.1.1.2	0	100	0	?	-	
N* i		100.1.1.7	0	100	0	?	-	

Total number of prefixes 21

Implementation Examples

Incomplete Global RPKI Adoption

Not all Internet Service Providers (ISPs) or regional networks fully participate in RPKI validation.

Example Scenario: Peers from regions with limited RPKI coverage advertise valid prefixes that appear as invalid due to missing ROAs.

Action: It may permit invalid routes from trusted peers or specific regions to maintain global connectivity.

Route Leak or Failover Scenarios

During failover or traffic engineering events, alternate paths may temporarily appear as invalid.

Example Scenario: A backup eBGP link originates a prefix from a different AS path than specified in the ROA.

Action: The route may be allowed conditionally (for example, through route maps) to ensure reachability during the transition period.

Commands

The BGP RPKI as origin validation feature introduces the following commands:

bgp rpki server

Use this command to configure an RPKI cache server using either TCP or SSH transport, and to set the port and timer parameters (refresh, retry, expire), along with authentication details when SSH is used.

Use the **no** parameter of this command to remove an existing RPKI server configuration from the BGP instance.

Command Syntax

```
bgp rpki server (A.B.C.D or X:X::X:X) (tcp|) (port (port number)) (refresh (1 - 86400 )) (retry (1 - 7200 )) (expire (600 - 17200))
```

```
bgp rpki server (A.B.C.D or X:X::X:X) ssh user (user name) encrypt (0|1) password (PASSWORD) (port (PORT NUMBER)) (refresh (1 - 86400 )) (retry (1 - 7200 )) (expire (600 - 17200))
```

Parameters

A.B.C.D

Specifies the IPv4 address of the RPKI cache server.

X:X::X:X

Specifies the IPv6 address of the RPKI cache server.

tcp

Specifies TCP as the communication protocol between the router and the RPKI server.

ssh

Specifies SSH as the communication protocol between the router and the RPKI server for secure communication.

user

Defines the username used to authenticate with the RPKI server.

encrypt

Specifies encrypted or not encrypted for a password

0

Defines unencrypted password (key)

1

Defines encrypted password (key)

password

Defines the BGP encrypted password (key) up to maximum 218 characters for an ssh connection.

port

Specifies the TCP or SSH port number to connect to the RPKI server.

refresh

Specifies the time interval, in seconds, to refresh the cache from the RPKI server.

retry

Specifies the time interval, in seconds, to retry the connection to the RPKI server if the previous attempt fails.

expire

Specifies the cache expiration interval, in seconds, after which cached RPKI data is discarded if not refreshed.

Default

None

Command Mode

Router mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to specify rpki server:

```
OcNOS(config-router)# bgp rpki server 10.30.0.85 tcp port 3323 refresh 600 retry 120 expire 7200  
OcNOS(config-router)# no bgp rpki server 10.30.0.85  
  
OcNOS(config-router)#bgp rpki server 1.1.1.1 ssh user test encrypt 0 password 123  
OcNOS(config-router)#commit
```

bgp origin-as validation-enable

Use this command to enable BGP Origin-AS (AS Origin) Validation using RPKI. When enabled, the router validates the origin AS of received BGP prefixes against the ROA information downloaded from RPKI servers..

Use the **no** parameter of this command to disable BGP Origin-AS validation.

Command Syntax

```
bgp origin-as validation-enable  
no bgp origin-as validation-enable
```

Parameters

None

Default

None

Command Mode

Address Family Configuration Mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example enables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config)#router bgp 100  
OcNOS(config-router)#address-family ipv4 unicast  
OcNOS(config-router-af)# bgp origin-as validation-enable
```

The following example disables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config-router-af)# no bgp origin-as validation-enable
```

bgp origin-as bestpath

Use this command to control how BGP selects the best path when RPKI Origin-AS validation is enabled.

This command determines whether BGP considers RPKI validation results during the bestpath selection process or allows paths with an invalid RPKI state to be selected.

Use the **no** form of this command to restore the default bestpath behavior.

Command Syntax

```
bgp origin-as bestpath (allow-invalid | use validity)
no bgp origin-as bestpath (allow-invalid | use validity)
```

Parameters

allow-invalid

Enables to handle a route with invalid RPKI state for the best path selection

Use-validity

Enables to use origin-as validation for the bestpath selection

Default

None

Command Mode

Address Family Configuration Mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example enables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config)#router bgp 100
OcNOS(config-router)#address-family ipv4 unicast
OcNOS(config-router-af)# bgp origin-as bestpath allow-invalid
```

The following example disables Origin-AS validation for BGP bestpath selection in the current address family configuration:

```
OcNOS(config-router-af)# no bgp origin-as bestpath allow-invalid
```


match rpki

Use this command to match BGP routes based on their Resource Public Key Infrastructure (RPKI) validation status in a route-map.

Use the **no** parameter of this command to remove an existing RPKI match configuration from the route-map.

Command Syntax

```
match rpki {valid | invalid | not-found}
no match rpki {valid | invalid | not-found}
```

Parameters

valid

Matches routes that have a valid RPKI validation status.

invalid

Matches routes that have an invalid RPKI validation status.

not-found

Matches routes whose RPKI validation status is unknown (not found).

Default

None

Command Mode

Route map mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates to match only invalid RPKI routes in a route-map and to remove the invalid RPKI routes.

```
ocnos#config terminal
ocnos(config)#route-map 1
ocnos(config-route-map)# match rpki invalid

ocnos(config-route-map)# no match rpki valid
```

show bgp rpki table ipv4

Use this command to display the IPv4 RPKI ROA table.

The command shows all validated ROA entries downloaded from configured RPKI servers.

Command Syntax

```
show bgp rpki table ipv4 (A.B.C.D/M ((covered|matched) (as-no <1-4294967295>)
```

Parameters

A.B.C.D/M

Displays ROA entries that include or relate to the specified IPv4 prefix.

covered

Displays ROA entries where the specified prefix is covered by a larger ROA prefix.

matched

Filters the ROA table to display entries associated with a specific authorized origin AS number.

as-no <1-4294967295>

Filters the ROA table to display entries associated with a specific authorized origin AS number.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays IPv4 RPKI ROA table:

```
ocnos#show bgp rpki table ipv4
BGP RPKI, ROA list
1.0.0.0/24                Maxlen:24 AS:13335 Server:10.30.0.85
                          Maxlen:24 AS:13335 Server:192.168.1.233
1.0.64.0/18              Maxlen:18 AS:18144 Server:10.30.0.85
                          Maxlen:18 AS:18144 Server:192.168.1.233
1.1.1.0/24               Maxlen:24 AS:13335 Server:10.30.0.85
                          Maxlen:24 AS:13335 Server:192.168.1.233
1.1.4.0/22               Maxlen:22 AS:4134 Server:10.30.0.85
                          Maxlen:22 AS:4134 Server:192.168.1.233
1.1.16.0/20              Maxlen:20 AS:4134 Server:10.30.0.85
                          Maxlen:20 AS:4134 Server:192.168.1.233
1.2.9.0/24               Maxlen:24 AS:4134 Server:10.30.0.85
                          Maxlen:24 AS:4134 Server:192.168.1.233
1.2.10.0/24              Maxlen:24 AS:4134 Server:10.30.0.85
                          Maxlen:24 AS:4134 Server:192.168.1.233
```

Explanation of output fields:

Field	Description
Prefix	Shows the IPv4 prefix (ROA prefix) published by the RPKI trust anchor.
Maxlen	Indicates the maximum prefix length allowed by the ROA.
AS	Displays the authorized Origin AS number for the prefix, as specified in the ROA.
Server	Indicates the RPKI server (RPKI RTR server address) from which the ROA entry was received.

show bgp rpki table ipv6

Use this command to display the IPv6 RPKI ROA table.

The command shows all validated ROA entries downloaded from configured RPKI servers.

Command Syntax

```
show bgp rpki table ipv6 (X:X::X:X/M ((covered|matched) (as-no <1-4294967295>|)|)|)
```

Parameters

X:X::X:X/M

Displays ROA entries that include or relate to the specified IPv6 prefix.

covered

Displays ROA entries where the specified prefix is covered by a larger ROA prefix.

matched

Filters the ROA table to display entries associated with a specific authorized origin AS number.

as-no <1-4294967295>

Filters the ROA table to display entries associated with a specific authorized origin AS number.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays IPv6 RPKI ROA table:

```
ocnos#show bgp rpki table ipv6
BGP RPKI, ROA list
2001:200::/32           Maxlen:32 AS:2500 Server:192.168.1.233
2001:200:136::/48       Maxlen:48 AS:9367 Server:192.168.1.233
2001:200:1ba::/48       Maxlen:48 AS:24047 Server:192.168.1.233
2001:200:900::/40       Maxlen:40 AS:7660 Server:192.168.1.233
2001:200:e00::/40       Maxlen:40 AS:4690 Server:192.168.1.233
2001:200:8000::/35      Maxlen:35 AS:4690 Server:192.168.1.233
2001:200:c000::/35      Maxlen:35 AS:23634 Server:192.168.1.233
2001:200:e000::/35      Maxlen:35 AS:7660 Server:192.168.1.233
2001:218::/32           Maxlen:32 AS:2914 Server:192.168.1.233
2001:218:2000:2::/64    Maxlen:64 AS:4058 Server:192.168.1.233
2001:218:2000:11::/64   Maxlen:64 AS:55569 Server:192.168.1.233
2001:218:2000:21::/64   Maxlen:64 AS:55569 Server:192.168.1.233
2001:218:2002::/48      Maxlen:48 AS:2914 Server:192.168.1.233
```

Explanation of output fields:

Field	Description
Prefix	Shows the IPv6 prefix (ROA prefix) published by the RPKI trust anchor.
Maxlen	Indicates the maximum prefix length allowed by the ROA.
AS	Displays the authorized Origin AS number for the prefix, as specified in the ROA.
Server	Indicates the RPKI server (RPKI RTR server address) from which the ROA entry was received.

show bgp origin-as validity ipv4

Use this command to display the RPKI Origin-AS validation state of IPv4 BGP routes. The command shows whether each route is classified as valid, not-found, or invalid based on the ROA information received from RPKI servers.

Command Syntax

```
show bgp origin-as validity ipv4 (valid|not-found|invalid|) (vrf WORD|)
```

Parameters

valid

Displays only the IPv4 prefixes that passed RPKI Origin-AS validation.

not-found

Displays routes for which no corresponding ROA entry exists.

invalid

Displays routes that fail RPKI Origin-AS validation.

vrf WORD

Displays RPKI validation results for the specified VRF instance.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays bgp origin-as validity for IPv4 address:

```
ocnos#show bgp origin-as validity ipv4
BGP table version is 9, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
                l - labeled, S Stale, x-EVPN MPLS
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

Origin-AS validation codes: V valid, I invalid, N not-found, D disabled

   Network        Next Hop        Metric    LocPrf    Weight Path    Ext-Color
V*>  1.6.136.0/24   100.2.1.3          0         100         0  300 9583 ?      -
N*>i  4.4.4.4/32     100.4.1.4          0         200         0  400 ?          -
N*    100.2.1.3     100.2.1.3          0         200         0  300 400 ?      -
I*>   5.5.5.5/32    100.2.1.3          0         300         0  300 9583 ?      -
N*>   100.3.1.0/24   100.2.1.3          0         200         0  300 9583 ?      -
I*>   100.4.1.0/24   100.2.1.3          0         300         0  300 9583 ?      -
N*>   172.16.181.0/24 100.2.1.3          0         200         0  300 9583 ?      -

Total number of prefixes 6
```

show bgp origin-as validity ipv6

Use this command to display the RPKI Origin-AS validation state of IPv6 BGP routes. The command shows whether each route is classified as valid, not-found, or invalid based on the ROA information received from RPKI servers.

Command Syntax

```
show bgp origin-as validity ipv6 (valid|not-found|invalid|) (vrf WORD|)
```

Parameters

valid

Displays only the IPv6 prefixes that passed RPKI Origin-AS validation.

not-found

Displays routes for which no corresponding ROA entry exists.

invalid

Displays routes that fail RPKI Origin-AS validation.

vrf WORD

Displays RPKI validation results for the specified VRF instance.

Default

None

Command Mode

Execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays origin-as validity for IPv6 address:

```
ocnos#show bgp origin-as validity ipv6
BGP table version is 270, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
                l - labeled, S Stale, x-EVPN MPLS
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

Origin-AS validation codes: V valid, I invalid, N not-found, D disabled

   Network          Next Hop                Metric   LocPrf   Weight Path
N*>i  2001::100:1:1:0/120
      2001::100:1:1:2 (fe80::20c:29ff:feea:6a30)
                                   0         100       0   ?       -
N*>   2001::100:2:1:0/120
      2001::100:2:1:3 (fe80::20c:29ff:fedf:6e11)
                                   0         100       0  300 ?       -
N*>   2001::100:3:1:0/120
      2001::100:2:1:3 (fe80::20c:29ff:fedf:6e11)
                                   0         100       0  300 ?       -
N* i   2001::100:4:1:4
      2001::100:4:1:4
                                   0         100       0  400 ?       -
N*>i  2001::100:4:1:0/120
```

	2001::100:1:1:2 (fe80::20c:29ff:feea:6a30)						
		0	100	0	?	-	
N*	2001::100:2:1:3 (fe80::20c:29ff:fedf:6e11)						
		0	100	0	300 400 ?	-	
N*>i	2601:647:6300:8dc0::/64						
	2001::100:1:1:2 (fe80::20c:29ff:feea:6a30)						
		0	100	0	?	-	
Total number of prefixes 5							

show bgp rpki server

Use this command to display the configured RPKI servers, their transport protocol, operational state, uptime, and the number of received Route Origin Authorizations (ROAs) for both IPv4 and IPv6.

Command Syntax

```
show bgp rpki server (detail|summary)
```

Parameters

None

Default

None

Command Mode

Privileged execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example displays rpki server show command:

```
ocnos#show bgp rpki server
Server Address          Transport  State      Uptime      ROAs (IPv4/IPv6)
10.16.99.108            TCP:3323  established 00:02:27    9/3
155.155.1.1             SSH:2222  established 00:02:26    9/3

ocnos#show bgp rpki server detail
BGP RPKI Server Information
Server Address: 155.155.1.1:
  Transport: TCP:3323
  RTR Version: 1
  State: established
  Synced: TRUE
  Uptime: 00:00:39
  ROAs (IPv4/IPv6): 9/3
  Configured Refresh-Interval: 15 seconds
  Configured Retry-Interval: 10 seconds
  Configured Expire-Interval: 600 seconds
  Actual Refresh-Interval: 5 seconds
  Actual Retry-Interval: 5 seconds
  Actual Expire-Interval: 600 seconds
  Rest of time to Refresh-Interval expiration: 2 seconds
  Rest of time to Expire-Interval expiration: 597 seconds
  ToBeDeleted: FALSE
```

Explanation of output fields:

Field	Description
Server Address	IP address of the configured RPKI server.
Transport	Transport protocol and port used to connect to the RPKI

Field	Description
	server (for example: TCP:or SSH:)
RTR Version	Version of the RPKI-to-Router (RTR) protocol negotiated with the server.
State	Current operational state of the RPKI server
Synced	Indicates whether the router has successfully synchronized ROA data with the server.
Uptime	Duration for which the server connection has been active or the current status since the last state change.
ROAs (IPv4/IPv6)	Number of valid Route Origin Authorizations received from the server, separated by IPv4 and IPv6 counts.
Configured Refresh-Interval	User-configured refresh timer for requesting updated validation data.
Configured Retry-Interval	User-configured retry interval for reconnecting after a session failure.
Configured Expire-Interval	User-configured maximum validity period for cached ROA data.
Actual Refresh-Interval	Refresh interval currently in use, as negotiated with the server.
Actual Retry-Interval	Retry interval currently in use, based on server negotiation.
Actual Expire-Interval	Expire interval currently in use, based on protocol negotiation.
Rest of time to Refresh-Interval expiration	Remaining time until the next refresh request is triggered.
Rest of time to Expire-Interval expiration	Remaining time before cached validation data becomes invalid.
ToBeDeleted	Indicates whether the server entry is marked for deletion.

Glossary

Key Terms/Acronym	Description
Border Gateway Protocol (BGP)	The standardized exterior gateway protocol used to exchange routing information between autonomous systems (ASes) on the Internet.
Resource Public Key Infrastructure (RPKI)	A framework designed to secure the Internet's routing infrastructure by cryptographically verifying that an AS is authorized to originate a specific IP prefix.
Origin Validation	A process where a BGP router validates the origin AS of a received route against RPKI data to determine if the route is legitimate.

Key Terms/Acronym	Description
Route Origin Authorization (ROA)	digitally signed object that specifies which AS is authorized to announce a particular IP prefix.

ECMP Support for L3EVPN

Overview

Equal Cost Multipath (ECMP) for L3EVPN helps in load-balancing the EVPN IRB traffic. It supports traffic on all the IP paths that are available in an MPLS based EVPN network with a symmetric IRB (S-IRB) configuration. This feature programs IP prefix advertisements received on multiple BGP paths into the forwarding plane, enabling ECMP load balancing of known unicast inter-subnet IP traffic.


Feature Characteristics

This feature supports the MPLS-based symmetric IRB (S-IRB) configuration between peer PEs to configure IP multipaths in forwarding-plane.

It applies specifically to MPLS-based EVPN deployments on Qumran1, Qumran2, and J2C+ platforms using VLAN-based or VLAN-aware bundle service interfaces between CE and PE. It ensures that hashing and L3 traffic load-balancing occurs at the ingress PE on a per-service level, independent of the transport-level hashing.

In the below image, the PE routers PE1 and PE2 learn IPv4 and IPv6 customer subnet routes from connected CE device into the corresponding EVPN MAC-VRF via VLAN-based or VLAN-aware bundle service interface. These PE devices' configurations have MAC-VRF connected to an IP-VRF via IRB interface. The protocol BGP running on these PE devices advertise the IRB IP/MAC, CE host IP/MAC, and customer subnet IP routes to UPE as Type-2 and Type-5 (carrying L3VNI for VXLAN-based EVPN or Layer-3 label for MPLS-based EVPN) routes containing the respective IP and MAC-VRF Route Distinguishers and Route Targets.

Figure 1. L3EVPN ECMP Characteristics



The User-facing PE (UPE) processes the advertised routes and their next-hops in its control-plane and installs them in its forwarding-plane as multipath route entries. It prepares the hashing keys based on the configured fields of the incoming traffic, and also load balances the traffic at EVPN service level (either Layer 2 or Layer 3) and independently at transport level (LDP/RSVP/SR-MPLS) .

BGP calculates the multipath routes at both RD and VRF levels and passes the route with multiple next-hops towards NSM. NSM and HSL use the existing ECMP infrastructure to install the EVPN routes with multipaths in fast

path. HSL uses hierarchical Forwarding Equivalence Class (FEC) and load balances the traffic at service level, within each ECMP group among its multipath member routes.

Benefits

Multipath programming of EVPN Type-2 and Type-5 BGP multipath routes into the data plane.

Operates over MPLS transport with support for LDP, RSVP-TE, or SR-MPLS.

L3/L4 load balancing of inter-subnet known unicast traffic at EVPN IRB service level.

Limitations

Support for symmetric IRB configuration only.

Compatible with P2P transport tunnels only.

Independent L2 and L3 load-balancing—this feature focuses only on L3 (IRB) ECMP.

Supported Hardware

Qumran1, Qumran2, and J2C+ devices.

Prerequisites

1. Define Interfaces and Loopback Addresses

Configure Layer 2 interfaces such as port channel interfaces, for example po1, and assign specific IP addresses for proper identification and routing. Additionally, assign loopback IP addresses to establish essential points of connectivity.

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 8.8.8.8/32 secondary  
  ipv6 address ::1/128  
  
interface po7  
  ip address 31.1.1.8/24
```

2. Configure IGP for Dynamic Routing

Enable ISIS to facilitate dynamic routing on all nodes within the network. Define ISIS router instances to match loopback IP addresses and add network segments to ISIS areas for proper route distribution. Set up neighbor relationships using loopback IP addresses, ensuring efficient route advertisement and convergence for optimal network performance.

ISIS Configuration

```
router isis 1  
  is-type level-2-only  
  metric-style wide  
  microloop-avoidance level-2  
  mpls traffic-eng router-id 8.8.8.8  
  mpls traffic-eng level-2  
  capability cspf  
  dynamic-hostname  
  bfd all-interfaces  
  net 49.0000.0000.0008.00  
  passive-interface lo  
!  
interface po7
```

```

mpls ldp-igp sync isis level-2
isis network point-to-point
ip router isis 1

```

Configuration

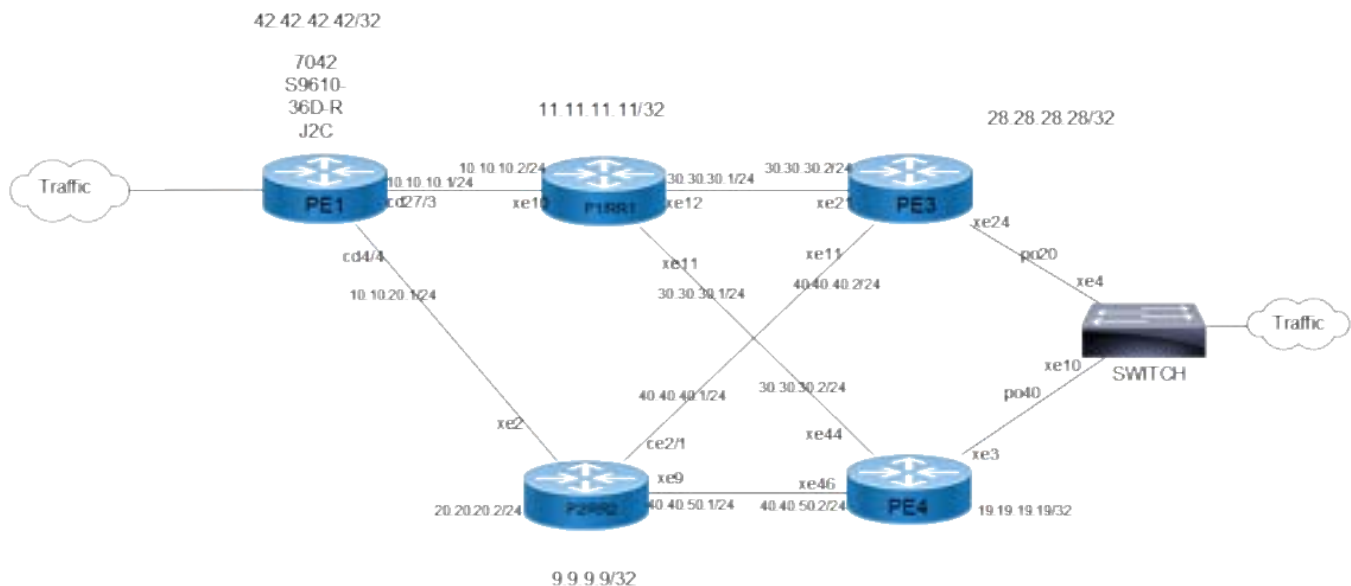
To set up ECMP for L3EVPN, follow the configuration steps mentioned below.

Topology

This topology shows the traffic distribution and load balancing across the equal cost paths from PE1 with EVPN service configured in symmetric IRB mode across all the PEs.

The LDP sessions are established between all P and PE routers.

Figure 2. L3EVPN ECMP Topology



PE1

1. Configure Router LDP

```

PE1(config)#router ldp
PE1(config-router)# router-id 8.8.8.8
PE1(config-router)# transport-address ipv4 8.8.8.8
PE1(config-router)#

```

2. Enable LDP and label-switching for core interface

```

PE1(config)#interface cd4
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

PE1(config)#interface cd27
PE1(config-if)# enable-ldp ipv4
PE1(config-if)#label-switching

```

3. Configure VRF

```
PE1(config)#ip vrf ip_vrfirb
PE1(config-vrf)# rd 90000:8
PE1(config-vrf)# route-target both 90000:1
PE1(config-vrf)#
```

4. Configure subinterface and attach VRF to subinterface

```
PE1(config)#interface xe26.100
PE1(config-if)# encapsulation dot1q 100
PE1(config-if)# ip vrf forwarding ip_vrfirb
PE1(config-if)# ip address 100.1.1.1/24
PE1(config-if)# ipv6 address 1000::1/24
```

5. Configure BGP

```
PE1(config)#router bgp 90000
PE1(config-router)# bgp router-id 8.8.8.8
PE1(config-router)# neighbor 4.4.4.4 remote-as 90000
PE1(config-router)# neighbor 4.4.4.4 update-source lo
PE1(config-router)# neighbor 7.7.7.7 remote-as 90000
PE1(config-router)# neighbor 7.7.7.7 update-source lo
PE1(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family ipv4 vrf ip_vrfirb
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# address-family ipv6 vrf ip_vrfirb
PE1(config-router-af)# redistribute connected
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE1(config-router)# exit
PE1(config)#
```

6. Configure Multipath and max-paths iBGP

```
PE1(config)#router bgp 90000
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#multi-path
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv4 vrf ip_vrfirb
PE1(config-router-af)#max-paths ibgp 4
PE1(config-router-af)#exit-address-family
PE1(config-router)#address-family ipv6 vrf ip_vrfirb
PE1(config-router-af)#max-paths ibgp 4
PE1(config-router-af)#
PE1(config-router-af)#exit
PE1(config-router)#
```

P1/RR1

1. Configure Router LDP

```
P1(config)#router ldp
P1(config-router)# router-id 4.4.4.4
P1(config-router)# transport-address ipv4 4.4.4.4
P1(config-router)#
```

2. Enable LDP and label-switching for core interface

```
P1(config)#interface xe12
P1(config-if)# enable-ldp ipv4
P1(config-if)#label-switching
```

```
P1(config)#interface xe10
P1(config-if)# enable-ldp ipv4
P1(config-if)#label-switching

P1(config)#interface xe11
P1(config-if)# enable-ldp ipv4
P1(config-if)#label-switching
```

3. Configure BGP

```
P1(config)#router bgp 90000
P1(config-router)# bgp router-id 4.4.4.4
P1(config-router)# no bgp inbound-route-filter
P1(config-router)# neighbor 3.3.3.3 remote-as 90000
P1(config-router)# neighbor 3.3.3.3 update-source lo
P1(config-router)# neighbor 5.5.5.5 remote-as 90000
P1(config-router)# neighbor 5.5.5.5 update-source lo
P1(config-router)# neighbor 8.8.8.8 remote-as 90000
P1(config-router)# neighbor 8.8.8.8 update-source lo
P1(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
P1(config-router)#
```

P2

1. Configure Router LDP

```
P2(config)#router ldp
P2(config-router)# router-id 7.7.7.7
P2(config-router)# transport-address ipv4 7.7.7.7
P2(config-router)#
```

2. Enable LDP and label-switching for core interface

```
P2(config)#interface xe2
P2(config-if)# enable-ldp ipv4
P2(config-if)#label-switching

P2(config)#interface xe9
P2(config-if)# enable-ldp ipv4
P2(config-if)#label-switching

P2(config)#interface ce2
P2(config-if)# enable-ldp ipv4
P2(config-if)#label-switching
```

3. Configure BGP

```
P2(config)#router bgp 90000
P2(config-router)# bgp router-id 7.7.7.7
P2(config-router)# no bgp inbound-route-filter
P2(config-router)# neighbor 3.3.3.3 remote-as 90000
P2(config-router)# neighbor 3.3.3.3 update-source lo
P2(config-router)# neighbor 5.5.5.5 remote-as 90000
P2(config-router)# neighbor 5.5.5.5 update-source lo
P2(config-router)# neighbor 8.8.8.8 remote-as 90000
P2(config-router)# neighbor 8.8.8.8 update-source lo
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
P2(config-router)#
```

PE3

1. Configure Router LDP

```
PE3(config)#router ldp
PE3(config-router)# router-id 5.5.5.5
PE3(config-router)# transport-address ipv4 5.5.5.5
PE3(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE3(config)#interface xe21
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

```
PE3(config)#interface xe24
PE3(config-if)# enable-ldp ipv4
PE3(config-if)#label-switching
```

3. Configure VRF

```
PE3(config)#ip vrf vrf1000
PE3(config-vrf)# rd 90000:5
PE3(config-vrf)# route-target both 90000:1
PE3(config-vrf)#
```

4. Configure sub-interface and attach the VRF to sub-interface

```
PE3(config)#interface xe2.100
PE3(config-if)# encapsulation dot1q 100
PE3(config-if)# ip vrf forwarding ip_vrfirb
PE3(config-if)# ip address 200.1.1.5/24
PE3(config-if)# ipv6 address 2000::5/64
PE3(config-if)#
```

5. Configure BGP

```
PE3(config)#router bgp 90000
PE3(config-router)# bgp router-id 5.5.5.5
PE3(config-router)# neighbor 4.4.4.4 remote-as 90000
PE3(config-router)# neighbor 4.4.4.4 update-source lo
PE3(config-router)# neighbor 7.7.7.7 remote-as 90000
PE3(config-router)# neighbor 7.7.7.7 update-source lo
PE3(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE3(config-router)# address-family ipv4 vrf ip_vrfirb
PE3(config-router-af)# redistribute connected
PE3(config-router-af)# exit-address-family
PE3(config-router)# !
PE3(config-router)# address-family ipv6 vrf ip_vrfirb
PE3(config-router-af)# redistribute connected
PE3(config-router-af)# exit-address-family
PE3(config-router)#
```

PE4

1. Configure Router LDP

```
PE4(config)#router ldp
PE4(config-router)# router-id 3.3.3.3
PE4(config-router)# transport-address ipv4 3.3.3.3
PE4(config-router)#
```

2. Enable LDP and label-switching for core interface

```
PE4(config)#interface xe3
```



```
PE4(config-if)# enable-ldp ipv4
PE4(config-if)#label-switching
PE4(config)#interface xe46
PE4(config-if)# enable-ldp ipv4
PE4(config-if)#label-switching
```

3. Configure VRF

```
PE4(config)#ip vrf vrf1000
PE4(config-vrf)# rd 90000:3
PE4(config-vrf)# route-target both 90000:1
PE4(config-vrf)#
```

4. Configure sub-interface and attach the VRF to sub-interface

```
PE4(config)#interface xe6.100
PE4(config-if)# encapsulation dot1q 100
PE4(config-if)# ip vrf forwarding ip_vrfirb
PE4(config-if)# ip address 200.1.1.3/24
PE4(config-if)# ipv6 address 2000::3/64
PE4(config-if)#
```

5. Configure BGP

```
PE4(config)#router bgp 90000
PE4(config-router)# bgp router-id 3.3.3.3
PE4(config-router)# neighbor 4.4.4.4 remote-as 90000
PE4(config-router)# neighbor 4.4.4.4 update-source lo
PE4(config-router)# neighbor 7.7.7.7 remote-as 90000
PE4(config-router)# neighbor 7.7.7.7 update-source lo
PE4(config-router)# !
PE1(config-router)#address-family l2vpn evpn
PE1(config-router-af)#neighbor PG activate
PE1(config-router-af)# exit-address-family
PE1(config-router)# !
PE4(config-router)# address-family ipv4 vrf ip_vrfirb
PE4(config-router-af)# redistribute connected
PE4(config-router-af)# exit-address-family
PE4(config-router)# !
PE4(config-router)# address-family ipv6 vrf ip_vrfirb
PE4(config-router-af)# redistribute connected
PE4(config-router-af)# exit-address-family
PE4(config-router)#
```

Configuration Snapshot

PE1

```
PE1-7042#sh running-config
!
! Software version: UFI_S9610-36D-OcNOS-SP-PLUS-7.0.0.136-Alpha 10/12/2025 17:35:02
!
! Last configuration change at 18:47:41 UTC Mon Oct 13 2025 by root
!
!
service password-encryption
!
no logging cli
logging logfile rsvp 7
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics voq-full-color enable
```

```
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE1-7042
port cd4 breakout 4X10g
port cd6 breakout 4X10g
port cd27 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn mpls irb
!
ip vrf ip_vrfirb
  rd 42.42.42.42:200
  route-target both 200:200
  l3vni 20000
!
ip vrf ip_vrfirbMH801
  rd 42.42.42.42:801
  route-target both 801:801
  l3vni 40000
!
ip vrf management
!
mac vrf vrfirb
  rd 42.42.42.42:2000
  route-target both 2000:2000
!
mac vrf vrfirbMH801
  rd 42.42.42.42:4000
  route-target both 4000:4000
!
evpn mpls vtep-ip-global 42.42.42.42
!
evpn mpls id 200
  host-reachability-protocol evpn-bgp vrfirb
!
evpn mpls id 801
  host-reachability-protocol evpn-bgp vrfirbMH801
  evpn irb irb700
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability node
  reoptimize
  exit
  inactivity-timer 5
  enable
  exit
!
```

```
interface cd0
!
interface cd1
!
interface cd2
!
interface cd3
!
interface cd4/1
!
interface cd4/2
!
interface cd4/3
!
interface cd4/4
load-interval 30
ip address 10.10.20.1/24
mtu 9216
label-switching
shutdown
ip router isis 1
enable-rsvp
!
interface cd5
!
interface cd6/1
!
interface cd6/2
!
interface cd6/3
!
interface cd6/4
!
interface cd6/4.200 switchport
encapsulation dot1q 200
rewrite pop
access-if-evpn
map vpn-id 200
!
interface cd6/4.801 switchport
encapsulation dot1q 801
rewrite pop
access-if-evpn
map vpn-id 801
!
load-balance enable
!
load-balance ipv4 src-dest-ipv4
load-balance ipv6 flow-label next-hdr src-dest-ipv6
load-balance src-dest-l4port
!
interface cd7
!
interface cd8
!
interface cd9
!
interface cd10
!
interface cd11
!
interface cd12
!
interface cd13
!
interface cd14
!
interface cd15
```

```
!  
interface cd16  
!  
interface cd17  
!  
interface cd18  
!  
interface cd19  
!  
interface cd20  
!  
interface cd21  
!  
interface cd22  
!  
interface cd23  
!  
interface cd24  
!  
interface cd25  
!  
interface cd26  
!  
interface cd27/1  
!  
interface cd27/2  
!  
interface cd27/3  
  load-interval 30  
  ip address 10.10.10.1/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-rsvp  
!  
interface cd27/4  
!  
interface cd28  
!  
interface cd29  
!  
interface cd30  
!  
interface cd31  
!  
interface cd32  
!  
interface cd33  
!  
interface cd34  
!  
interface cd35  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface irb100  
  ip vrf forwarding ip_vrfirb  
  ip address 80.80.1.1/24 anycast  
  ipv6 address 80::1/64 anycast  
!  
interface irb700  
  ip vrf forwarding ip_vrfirbMH801  
  ip address 81.81.1.1/24 anycast  
  ipv6 address 81::1/64 anycast  
!  
interface lo
```

```
ip address 127.0.0.1/8
ip address 42.42.42.42/32 secondary
ipv6 address ::1/128
ip router isis 1
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
evpn irb-forwarding anycast-gateway-mac 0011.2200.0001
!
exit
!
router ospf
!
router ospf 100
bfd all-interfaces
network 10.10.10.0/24 area 0.0.0.0
network 10.10.20.0/24 area 0.0.0.0
network 42.42.42.42/32 area 0.0.0.0
!
router bgp 65100
bgp router-id 42.42.42.42
no bgp inbound-route-filter
allocate-label all
neighbor 11.11.11.11 remote-as 65100
neighbor 11.11.11.11 update-source lo
neighbor 81.81.1.2 remote-as 65101
!
address-family ipv4 unicast
network 42.42.42.42/32
exit-address-family
!
address-family l2vpn evpn
multi-path
neighbor 11.11.11.11 activate
neighbor 11.11.11.11 route-reflector-client
exit-address-family
!
address-family ipv4 vrf ip_vrfirb
redistribute connected
exit-address-family
!
address-family ipv4 vrf ip_vrfirbMH801
max-paths ibgp 4
redistribute connected
neighbor 81.81.1.2 remote-as 65101
neighbor 81.81.1.2 activate
exit-address-family
!
exit
!
rsvp-path PE1-PE3 mpls
10.10.10.2 strict
30.30.40.2 strict
!
rsvp-path PE1-PE4 mpls
10.10.10.2 strict
30.30.30.2 strict
!
rsvp-trunk PE1-PE2 ipv4
to 30.30.30.30
!
rsvp-trunk PE1-PE3 ipv4
no primary cspf
primary path PE1-PE3
to 28.28.28.28
```

```

!
rsvp-trunk PE1-PE4 ipv4
  no primary cspf
  primary path PE1-PE4
  to 19.19.19.19
!
!
end

!
PE1-7042#

PE1-7042#sh ip ospf neighbor

Total number of full neighbors: 0
OSPF process 0 VRF(default):
Neighbor ID      Pri   State                Dead Time   Address      Interface     Instance ID

Total number of full neighbors: 1
OSPF process 100 VRF(default):
Neighbor ID      Pri   State                Dead Time   Address      Interface     Instance ID
11.11.11.11      1     Full/Backup          00:00:34    10.10.10.2    cd27/3        0
PE1-7042#
PE1-7042#sh rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To      From      Tun-ID   LSP-ID   Type   LSPName      State
Uptime  Rt  Style  Labelin  Labelout
30.30.30.30  42.42.42.42  5001     2201     PRI    PE1-PE2-
Primary                                UP      00:20:33  1 1 SE    -      24327
28.28.28.28  42.42.42.42  5002     2203     PRI    PE1-PE3-
Primary                                UP      00:21:28  1 1 SE    -      24326
19.19.19.19  42.42.42.42  5003     2204     PRI    PE1-PE4-
Primary                                UP      00:21:28  1 1 SE    -      24325
Total 3 displayed, Up 3, Down 0.

Egress RSVP:
To      From      Tun-ID   LSP-ID   Type   LSPName      State
Uptime  Rt  Style  Labelin  Labelout
42.42.42.42  19.19.19.19  5001     2201     PRI    PE4-PE1-
Primary                                UP      00:20:26  1 1 SE    24960   -
42.42.42.42  28.28.28.28  5001     2201     PRI    PE3-PE1-
Primary                                UP      00:20:01  1 1 SE    24962   -
42.42.42.42  30.30.30.30  5001     2202     PRI    PE2-PE1-
Primary                                UP      00:20:21  1 1 SE    24961   -
Total 3 displayed, Up 3, Down 0.

PE1-7042#
PE1-7042#
PE1-7042#show bgp l2vpn evpn
BGP table version is 4, local router ID is 42.42.42.42
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route

```

5 - Prefix Route

Network	Next Hop	Metric	LocPrf	Weight	Path	Peer	Encap
RD[65100:501]							
*>i [3]:[501]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:502]							
*>i [3]:[502]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:503]							
*>i [3]:[503]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:504]							
*>i [3]:[504]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:505]							
*>i [3]:[505]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:506]							
*>i [3]:[506]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:507]							
*>i [3]:[507]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:508]							
*>i [3]:[508]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:509]							
*>i [3]:[509]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:510]							
*>i [3]:[510]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[19.19.19.19:250]							
*>i [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[16]	19.19.19.19	0	100	0 ?	-	11.11.11.11	MPLS
RD[19.19.19.19:801]							
*>i [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]	19.19.19.19	0	100	0 ?	-	11.11.11.11	MPLS
RD[19.19.19.19:3000]							
*>i [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]	19.19.19.19	0	100	0 i	-	11.11.11.11	MPLS
*>i [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]	19.19.19.19	0	100	0 i	-	11.11.11.11	MPLS
*>i [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]	19.19.19.19	0	100	0 i	-	11.11.11.11	MPLS
*>i [3]:[251]:[32,19.19.19.19]	19.19.19.19	0	100	0 i	-	11.11.11.11	MPLS
RD[19.19.19.19:4000]							
*>i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]	19.19.19.19	0	100	0 i	-	11.11.11.11	MPLS
*>i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]	19.19.19.19	0	100	0 i	-	11.11.11.11	MPLS
*>i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]							

```

19.19.19.19      0      100      0      i      -      11.11.11.11      MPLS
*>i      [3]:[811]:[32,19.19.19.19]
19.19.19.19      0      100      0      i      -      11.11.11.11      MPLS

RD[19.19.19.19:64512]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
19.19.19.19      0      100      0      i      -      11.11.11.11      MPLS
*>i      [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
19.19.19.19      0      100      0      i      -      11.11.11.11      MPLS

RD[28.28.28.28:200]
*>i      [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
28.28.28.28      0      100      0      ?      -      11.11.11.11      MPLS

RD[28.28.28.28:801]
*>i      [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
28.28.28.28      0      100      0      ?      -      11.11.11.11      MPLS

RD[28.28.28.28:2000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [3]:[201]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS

RD[28.28.28.28:4000]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [2]:[00:00:00:11:11:22:22:00:00:00]:[811]:[48,0010:9400:0005]:[0]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [3]:[811]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS

RD[28.28.28.28:64512]
*>i      [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>i      [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS

RD[30.30.30.30:250]
*>i      [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
30.30.30.30      0      100      0      ?      -      11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i      [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
30.30.30.30      0      100      0      i      -      11.11.11.11      MPLS
*>i      [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
30.30.30.30      0      100      0      i      -      11.11.11.11      MPLS
*>i      [3]:[250]:[32,30.30.30.30]
30.30.30.30      0      100      0      i      -      11.11.11.11      MPLS

RD[42.42.42.42:2000] VRF[vrfirb]
* i      [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
* i      [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
28.28.28.28      0      100      0      i      -      11.11.11.11      MPLS
*>      [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
42.42.42.42      0      100      32768      i      -      -----
MPLS
*>      [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
42.42.42.42      0      100      32768      i      -      -----
MPLS

```



```

*> [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0008]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0009]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000a]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000b]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000c]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000d]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000e]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:000f]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0010]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0011]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0012]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0013]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0014]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0015]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0016]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0017]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0018]:[0]:[18]
      42.42.42.42      0      100      32768      i      -      -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0019]:[0]:[18]

```

```

42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001a]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001b]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001c]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001d]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001e]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:001f]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0020]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0021]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0022]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0023]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0024]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0025]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0026]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0027]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0028]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0029]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002a]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002b]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002c]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002d]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002e]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:002f]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0030]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS

```

```

*> [2]:[0]:[200]:[48,1200:0100:0031]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0032]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0033]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0034]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0035]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0036]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0037]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0038]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0039]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003a]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003b]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003c]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003d]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003e]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:003f]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0040]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0041]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0042]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0043]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0044]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0045]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0046]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0047]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0048]:[0]:[18]

```

```

42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0049]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004a]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004b]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004c]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004d]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004e]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:004f]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0050]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0051]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0052]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0053]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
42.42.42.42      0      100      32768  i      -  -----
MPLS

```

```

*> [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
      42.42.42.42      0      100      32768 i      - -----
MPLS
* i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1][19]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
*> [3]:[200]:[32,42.42.42.42]
      42.42.42.42      0      100      32768 i      - -----
MPLS
* i [3]:[201]:[32,28.28.28.28]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS

RD[42.42.42.42:4000] VRF[vrfirbMH801]
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i      19.19.19.19      0      100      0 i      - 11.11.11.11 MPLS
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i      19.19.19.19      0      100      0 i      - 11.11.11.11 MPLS
*> [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1][19]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42      0      100      32768 i      - -----
MPLS
*> [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1][19]
      42.42.42.42      0      100      32768 i      - -----
MPLS
* i [2]:[00:00:00:11:11:22:22:00:00:00]:[811]:[48,0010:9400:0005]:[0]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i      19.19.19.19      0      100      0 i      - 11.11.11.11 MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1][20]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS
* i      19.19.19.19      0      100      0 i      - 11.11.11.11 MPLS
*> [3]:[801]:[32,42.42.42.42]
      42.42.42.42      0      100      32768 i      - -----
MPLS
* i [3]:[811]:[32,19.19.19.19]
      19.19.19.19      0      100      0 i      - 11.11.11.11 MPLS
* i [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      0 i      - 11.11.11.11 MPLS

Total number of prefixes 158
PE1-7042#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

```

```

Code   FEC
Label  Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID  Nhlfe-ID  Pri  Out-
B(x)>   90.90.1.0/24      00:20:35 -      1      2      13      -   -
      -                  -      -                  -      12      Yes  16   -
      28.28.28.28      -      -                  -
B(x)>   91.91.1.0/24      00:20:35 -      1      3      15      (e)  -   -
      -                  -      -                  -
      19.19.19.19      -      -                  -
      28.28.28.28      -      -                  -
PE1-7042#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0
IP Route Table for VRF "ip_vrfirb"
B      > 90.90.1.0/24 [200/0] via 28.28.28.28, installed 00:20:35, last update 00:20:35 ago

Total number of IPv4 routes 1
IP Route Table for VRF "ip_vrfirbMH801"
B      > 91.91.1.0/24 [200/0] via 28.28.28.28, installed 00:20:35, last update 00:20:35 ago
      > [200/0] via 19.19.19.19

Total number of IPv4 routes 2
IP Route Table for VRF "vrfirb"

Total number of IPv4 routes 0
IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0

Gateway of last resort is not set
PE1-7042#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF | DESTINATION | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS | OUT | OUT | NEXTHOP | | | |
| ID | FEC | cnt | cnt | Ifname | IFNAME | LABEL |
| OBJECT | | | | | | |
+-----+-----+-----+-----+-----+-----+-----+-----+
2  B> 90.90.1.0/24      13(s) -      -      -      -
      0x2000ccdb 00:22:34 1 1
      - /- /PRI /12 -
      16 28.28.28.28 0x2000ccdb
3  B> 91.91.1.0/24      15(e) -      -      -      -
      0x20000001 00:22:34 1 1
      - /- /PRI /10 -
      17 19.19.19.19 0x2000ccdc
      - /- /PRI /14 -
      17 28.28.28.28 0x2000ccdd

PE1-7042#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color FAT
=====
42.42.42.42 19.19.19.19      Installed   00:22:37     00:22:37     40000      ---
---
42.42.42.42 28.28.28.28      Installed   00:22:37     00:22:37     40000      ---
---
42.42.42.42 28.28.28.28      Installed   00:22:37     00:22:37     20000      ---

```

```

      ---      ---      ---

Total number of entries are 3
PE1-7042#

apps-fileview.texmex_20251030.00_p0
PE1-7042.txt
Displaying PE1-7042.txt.

```

P1/RR1

```

P1-7011#sh running-config
!
! Software version: HFCL-AR-OcNOS-CSR-7.0.0.133-Alpha
10/09/20                                     25 17:37:30
!
! Last configuration change at 11:38:04 UTC Sat Aug 02 2025
by                                           root
!
!
service password-encryption
!
logging console 5
logging monitor disable
logging logfile rsvp 7
logging level nsm 4
logging level rsvp 7
logging level bgp 7
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname P1-7011
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
ip vrf management
!
router ldp
!
router rsvp
!
interface eth0
  ip vrf forwarding management

```

```
ip address dhcp
!
interface ge1
!
interface ge2
!
interface ge3
!
interface ge4
!
interface lo
ip address 127.0.0.1/8
ip address 11.11.11.11/32 secondary
ipv6 address ::1/128
ip router isis 1
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe5
!
interface xe6
!
interface xe7
load-interval 30
ip address 20.20.30.2/24
mtu 9216
label-switching
ip router isis 1
enable-ldp ipv4
enable-rsvp
!
interface xe8
!
interface xe9
!
interface xe10
load-interval 30
ip address 10.10.10.2/24
mtu 9216
label-switching
ip router isis 1
enable-ldp ipv4
enable-rsvp
!
interface xe11
load-interval 30
ip address 30.30.30.1/24
mtu 9216
label-switching
ip router isis 1
enable-ldp ipv4
enable-rsvp
!
interface xe12
load-interval 30
ip address 30.30.40.1/24
mtu 9216
label-switching
ip router isis 1
enable-ldp ipv4
enable-rsvp
!
exit
!
router ospf 100
```



```

bfd all-interfaces
network 10.10.10.0/24 area 0.0.0.0
network 11.11.11.11/32 area 0.0.0.0
network 20.20.30.0/24 area 0.0.0.0
network 30.30.30.0/24 area 0.0.0.0
network 30.30.40.0/24 area 0.0.0.0
!
router isis 1
 is-type level-2-only
 metric-style wide
 microloop-avoidance level-2
 mpls traffic-eng router-id 11.11.11.11
 mpls traffic-eng level-2
 dynamic-hostname
 fast-reroute per-prefix level-2 proto ipv4 all
 fast-reroute ti-lfa level-2 proto ipv4
 bfd all-interfaces
 net 49.0000.0000.0011.00
!
router bgp 65100
 bgp router-id 11.11.11.11
 no bgp inbound-route-filter
 allocate-label all
 neighbor 19.19.19.19 remote-as 65100
 neighbor 19.19.19.19 update-source lo
 neighbor 28.28.28.28 remote-as 65100
 neighbor 28.28.28.28 update-source lo
 neighbor 30.30.30.30 remote-as 65100
 neighbor 30.30.30.30 update-source lo
 neighbor 42.42.42.42 remote-as 65100
 neighbor 42.42.42.42 update-source lo
!
 address-family l2vpn evpn
 neighbor 19.19.19.19 activate
 neighbor 19.19.19.19 route-reflector-client
 neighbor 28.28.28.28 activate
 neighbor 28.28.28.28 route-reflector-client
 neighbor 30.30.30.30 activate
 neighbor 30.30.30.30 route-reflector-client
 neighbor 42.42.42.42 activate
 neighbor 42.42.42.42 route-reflector-client
 exit-address-family
!
exit
!
!
end

!
P1-7011#

P1-7011#sh ip ospf neighbor

Total number of full neighbors: 4
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address        Interface      Instance ID
42.42.42.42      1     Full/DR         00:00:35   10.10.10.1     xe10           0
30.30.30.30      1     Full/DR         00:00:35   20.20.30.1     xe7            0
19.19.19.19      1     Full/DR         00:00:38   30.30.30.2     xe11           0
28.28.28.28      1     Full/DR         00:00:39   30.30.40.2     xe12           0
P1-7011#
P1-7011#sh rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

```

Transit RSVP:

To	Uptime	Rt	From	Style	Labelin	Tun-ID	Labelout	LSP-ID	Type	LSPName	State
30.30.30.30			42.42.42.42			5001		2201	PRI	PE1-PE2-	
Primary			UP			00:25:19		1 1 SE	24327	24962	
42.42.42.42			19.19.19.19			5001		2201	PRI	PE4-PE1-	
Primary			UP			00:25:12		1 1 SE	24328	24960	
42.42.42.42			28.28.28.28			5001		2201	PRI	PE3-PE1-	
Primary			UP			00:24:47		1 1 SE	24330	24962	
42.42.42.42			30.30.30.30			5001		2202	PRI	PE2-PE1-	
Primary			UP			00:25:07		1 1 SE	24329	24961	
28.28.28.28			30.30.30.30			5002		2201	PRI	PE2-PE3-	
Primary			UP			00:29:36		1 1 SE	24321	24960	
28.28.28.28			42.42.42.42			5002		2203	PRI	PE1-PE3-	
Primary			UP			00:26:14		1 1 SE	24326	24962	
30.30.30.30			19.19.19.19			5002		2202	PRI	PE4-PE2-	
Primary			UP			00:27:01		1 1 SE	24323	24961	
30.30.30.30			28.28.28.28			5002		2202	PRI	PE3-PE2-	
Primary			UP			00:29:42		1 1 SE	24320	24960	
19.19.19.19			28.28.28.28			5003		2203	PRI	PE3-PE4-	
Primary			UP			00:26:46		1 1 SE	24324	24960	
19.19.19.19			42.42.42.42			5003		2204	PRI	PE1-PE4-	
Primary			UP			00:26:14		1 1 SE	24325	24961	
28.28.28.28			19.19.19.19			5003		2203	PRI	PE4-PE3-	
Primary			UP			00:27:10		1 1 SE	24322	24961	

Total 11 displayed, Up 11, Down 0.

P1-7011#

P1-7011#

P1-7011#show bgp l2vpn evpn

BGP table version is 10, local router ID is 11.11.11.11

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,

l - labeled, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]

1 - Ethernet Auto-discovery Route

2 - MAC/IP Route

3 - Inclusive Multicast Route

4 - Ethernet Segment Route

5 - Prefix Route

Network	Next Hop	Metric	LocPrf	Weight	Path	Peer	Encap
RD[65100:501]							
*>i [3]:[501]:[32,30.30.30.30]	30.30.30.30	0	100	0	i	- 30.30.30.30	MPLS
RD[65100:502]							
*>i [3]:[502]:[32,30.30.30.30]	30.30.30.30	0	100	0	i	- 30.30.30.30	MPLS
RD[65100:503]							
*>i [3]:[503]:[32,30.30.30.30]	30.30.30.30	0	100	0	i	- 30.30.30.30	MPLS
RD[65100:504]							
*>i [3]:[504]:[32,30.30.30.30]	30.30.30.30	0	100	0	i	- 30.30.30.30	MPLS
RD[65100:505]							
*>i [3]:[505]:[32,30.30.30.30]	30.30.30.30	0	100	0	i	- 30.30.30.30	MPLS
RD[65100:506]							
*>i [3]:[506]:[32,30.30.30.30]	30.30.30.30	0	100	0	i	- 30.30.30.30	MPLS

```

RD[65100:507]
*>i    [3]:[507]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          - 30.30.30.30    MPLS

RD[65100:508]
*>i    [3]:[508]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          - 30.30.30.30    MPLS

RD[65100:509]
*>i    [3]:[509]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          - 30.30.30.30    MPLS

RD[65100:510]
*>i    [3]:[510]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          - 30.30.30.30    MPLS

RD[19.19.19.19:250]
*>i    [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[16]
          19.19.19.19          0          100          0    ?          - 19.19.19.19    MPLS

RD[19.19.19.19:801]
*>i    [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
          19.19.19.19          0          100          0    ?          - 19.19.19.19    MPLS

RD[19.19.19.19:3000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [3]:[251]:[32,19.19.19.19]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS

RD[19.19.19.19:4000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [3]:[811]:[32,19.19.19.19]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS

RD[19.19.19.19:64512]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS
*>i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
          19.19.19.19          0          100          0    i          - 19.19.19.19    MPLS

RD[28.28.28.28:200]
*>i    [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
          28.28.28.28          0          100          0    ?          - 28.28.28.28    MPLS

RD[28.28.28.28:801]
*>i    [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
          28.28.28.28          0          100          0    ?          - 28.28.28.28    MPLS

RD[28.28.28.28:2000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
          28.28.28.28          0          100          0    i          - 28.28.28.28    MPLS
*>i    [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
          28.28.28.28          0          100          0    i          - 28.28.28.28    MPLS
*>i    [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
          28.28.28.28          0          100          0    i          - 28.28.28.28    MPLS
*>i    [3]:[201]:[32,28.28.28.28]
          28.28.28.28          0          100          0    i          - 28.28.28.28    MPLS

```

```

RD[28.28.28.28:4000]
*>i   [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0      i      - 28.28.28.28      MPLS
*>i   [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0      i      - 28.28.28.28      MPLS
*>i   [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      28.28.28.28      0      100      0      i      - 28.28.28.28      MPLS
*>i   [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 28.28.28.28      MPLS

RD[28.28.28.28:64512]
*>i   [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0      i      - 28.28.28.28      MPLS
*>i   [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 28.28.28.28      MPLS

RD[30.30.30.30:250]
*>i   [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0      100      0      ?      - 30.30.30.30      MPLS

RD[30.30.30.30:3000]
*>i   [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0      100      0      i      - 30.30.30.30      MPLS
*>i   [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
      30.30.30.30      0      100      0      i      - 30.30.30.30      MPLS
*>i   [3]:[250]:[32,30.30.30.30]
      30.30.30.30      0      100      0      i      - 30.30.30.30      MPLS

RD[42.42.42.42:801]
*>i   [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
      42.42.42.42      0      100      0      ?      - 42.42.42.42      MPLS

RD[42.42.42.42:2000]
*>i   [3]:[200]:[32,42.42.42.42]
      42.42.42.42      0      100      0      i      - 42.42.42.42      MPLS

RD[42.42.42.42:4000]
*>i   [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42      0      100      0      i      - 42.42.42.42      MPLS
*>i   [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
      42.42.42.42      0      100      0      i      - 42.42.42.42      MPLS
*>i   [3]:[801]:[32,42.42.42.42]
      42.42.42.42      0      100      0      i      - 42.42.42.42      MPLS

Total number of prefixes 43
Pl-7011#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC
Label  Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID      Nhlfe-ID      Pri      Out-
Label  Out-Intf      Nexthop      UpTime      Ext-Color

Pl-7011#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0

Gateway of last resort is not set
Pl-7011#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup

```

L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy

```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF |      DESTINATION      | TNL/LSP/TYPE/NHLFE-IX |
|      MPLS      | OUT | OUT | NEXTHOP | EGRESS | UpTime | Ref |Prefix|
|
ID |      FEC      | Ifname | IFNAME | LABEL |
|      OBJECT | cnt | cnt |
+-----+-----+-----+-----+-----+-----+-----+-----+
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

```

P1-7011#sh evpn mpls tunnel
% Evpn mpls is not Enabled !
P1-7011#

```

```

apps-fileview.texmex_20251030.00_p0
PlRR.txt
Displaying PlRR.txt.

```

P2/RR2

```

P2-7009#sh running-config
!
! Software version: HFCL-CUAR-OcNOS-SP-PLUS-7.0.0.133-Alpha
10 /09/2025 17:37:22
!
! Last configuration change at 14:33:47 UTC Sun Sep 07 2025
by root
!
!
service password-encryption
!
logging console 5
logging logfile rsvp 7
logging level nsm 4
logging level rsvp 7
logging level bgp 7
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
hostname P2-7009
port ce2 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
feature rsyslog
logging remote server 10.16.58.70 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!

```

```
ip vrf management
!
router ldp
!
router rsvp
!
interface ce1
!
interface ce2/1
  load-interval 30
  ip address 40.40.40.1/24
  mtu 9216
  label-switching
  ip router isis 1
  enable-ldp ipv4
  enable-rsvp
!
interface ce2/2
!
interface ce2/3
!
interface ce2/4
!
interface ce3
!
interface ce4
!
interface ce5
!
interface ce6
!
interface ce7
!
interface ce8
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 9.9.9.9/32 secondary
  ipv6 address ::1/128
  ip router isis 1
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xe1
!
interface xe2
  speed 10g
  load-interval 30
  ip address 10.10.20.2/24
  mtu 9216
  label-switching
  ip router isis 1
  enable-ldp ipv4
  enable-rsvp
!
interface xe3
!
interface xe4
!
interface xe5
  load-interval 30
```

```
ip address 20.20.20.2/24
mtu 9216
label-switching
ip router isis 1
enable-ldp ipv4
enable-rsvp
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
speed 10g
load-interval 30
ip address 40.40.50.1/24
mtu 9216
label-switching
ip router isis 1
enable-ldp ipv4
enable-rsvp
!
interface xe10
!
interface xe11
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
exit
!
router ospf 100
bfd all-interfaces
network 9.9.9.9/32 area 0.0.0.0
network 10.10.20.0/24 area 0.0.0.0
network 20.20.20.0/24 area 0.0.0.0
network 40.40.40.0/24 area 0.0.0.0
network 40.40.50.0/24 area 0.0.0.0
!
router isis 1
is-type level-2-only
metric-style wide
microloop-avoidance level-2
mpls traffic-eng router-id 9.9.9.9
mpls traffic-eng level-2
dynamic-hostname
fast-reroute per-prefix level-2 proto ipv4 all
fast-reroute ti-lfa level-2 proto ipv4
bfd all-interfaces
net 49.0000.0000.0009.00
!
!
```

```

end

!
P2-7009#

P2-7009#sh ip ospf neighbor

Total number of full neighbors: 3
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address      Interface     Instance ID
30.30.30.30      1    Full/DR         00:00:32   20.20.20.1   xe5           0
28.28.28.28      1    Full/DR         00:00:35   40.40.40.2   ce2/1         0
19.19.19.19      1    Full/DR         00:00:32   40.40.50.2   xe9           0
P2-7009#
P2-7009#sh rsvp session
Type  : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

P2-7009#
P2-7009#
P2-7009#show bgp l2vpn evpn

Total number of prefixes 0
P2-7009#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code   FEC
Label   Out-Intf      Nexthop      UpTime      FTN-ID VRF-ID   Nhlfe-ID   Pri   Out-
P2-7009#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0

Gateway of last resort is not set
P2-7009#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF |          DESTINATION          | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS | OUT | OUT | NEXTHOP |                                     |                                     |                                     |
| ID |          FEC          | cnt | cnt | Ifname | IFNAME | LABEL |
| OBJECT |                                     |                                     |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+
P2-7009#sh evpn mpls tunnel
% Evpn mpls is not Enabled !
P2-7009#

apps-fileview.texmex_20251030.00_p0
P2.txt
Displaying P2.txt.

```


PE3

```
PE3-7028#sh running-config
!
! Software version: HFCL-DUAR-OcNOS-SP-PLUS-7.0.0.136-Alpha
10/1
!
! Last configuration change at 13:45:22 UTC Fri Sep 12 2025 by
r
!
!
service password-encryption
!
logging console 5
logging logfile rsvp 7
logging level nsm 4
logging level rsvp 7
logging level bgp 7
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter evpn-mpls-mh enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
qos enable
!
hostname PE3-7028
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn mpls irb
!
evpn mpls multihoming enable
!
ip vrf ip_vrfirb
  rd 28.28.28.28:200
  route-target both 200:200
  l3vni 20000
!
ip vrf ip_vrfirbMH801
  rd 28.28.28.28:801
  route-target both 801:801
  l3vni 40000
!
ip vrf management
!
mac vrf vrfirb
  rd 28.28.28.28:2000
```

```
    route-target both 2000:2000
!
mac vrf vrfirbMH801
  rd 28.28.28.28:4000
  route-target both 4000:4000
!
evpn mpls vtep-ip-global 28.28.28.28
!
evpn mpls id 102 xconnect target-mpls-id 101
!
evpn mpls id 201
  host-reachability-protocol evpn-bgp vrfirb
  evpn irb irb101
!
evpn mpls id 811
  host-reachability-protocol evpn-bgp vrfirbMH801
  evpn irb irb711
!
router rsvp
  auto-bypass
  attributes best-effort
  protection-capability node
  reoptimize
  exit
  inactivity-timer 5
  enable
  exit
!
interface po20
  switchport
  evpn multi-homed system-mac 0000.1111.2222
!
interface po20.201 switchport
  encapsulation dot1q 201
  rewrite pop
  access-if-evpn
  map vpn-id 201
!
interface po20.811 switchport
  encapsulation dot1q 811
  rewrite pop
  access-if-evpn
  map vpn-id 811
!
interface cel
!
interface ce2
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface ge25
!
interface irb101
  ip vrf forwarding ip_vrfirb
  ip address 90.90.1.1/24 anycast
  ipv6 address 90::1/64 anycast
!
interface irb711
  ip vrf forwarding ip_vrfirbMH801
  evpn irb-if-forwarding anycast-gateway-mac
  ip address 91.91.1.1/24 anycast
  ipv6 address 91::1/64 anycast
  ip ospf network point-to-point
!
interface lo
  ip address 127.0.0.1/8
```

```
ip address 28.28.28.28/32 secondary
ipv6 address ::1/128
ip router isis 1
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe1
!
interface xe2
!
interface xe3
!
interface xe4
!
interface xe5
!
interface xe6
!
interface xe7
!
interface xe8
!
interface xe9
!
interface xe10
!
interface xe11
load-interval 30
ip address 40.40.40.2/24
mtu 9216
label-switching
ip router isis 1
enable-rsvp
!
interface xe12
!
interface xe13
!
interface xe14
!
interface xe15
!
interface xe16
!
interface xe17
!
interface xe18
!
interface xe19
!
interface xe20
!
interface xe21
load-interval 30
ip address 30.30.40.2/24
mtu 9216
label-switching
ip router isis 1
enable-rsvp
!
interface xe22
!
interface xe23
!
interface xe24
```

```
speed 10g
channel-group 20 mode active
!
evpn irb-forwarding anycast-gateway-mac 0011.2200.0002
!
exit
!
router ospf 100
bfd all-interfaces
network 28.28.28.28/32 area 0.0.0.0
network 30.30.40.0/24 area 0.0.0.0
network 40.40.40.0/24 area 0.0.0.0
!
router bgp 65100
bgp router-id 28.28.28.28
no bgp inbound-route-filter
allocate-label all
neighbor 11.11.11.11 remote-as 65100
neighbor 11.11.11.11 update-source lo
neighbor 91.81.1.2 remote-as 65102
!
address-family ipv4 unicast
network 28.28.28.28/32
exit-address-family
!
address-family l2vpn evpn
neighbor 11.11.11.11 activate
neighbor 11.11.11.11 route-reflector-client
exit-address-family
!
address-family ipv4 vrf ip_vrfirb
redistribute connected
exit-address-family
!
address-family ipv4 vrf ip_vrfirbMH801
max-paths ibgp 4
redistribute connected
redistribute static
neighbor 91.91.1.2 remote-as 65102
neighbor 91.91.1.2 activate
exit-address-family
!
exit
!
rsvp-path PE3-PE1 mpls
30.30.40.1 strict
10.10.10.1 strict
!
rsvp-trunk PE3-PE1 ipv4
no primary cspf
primary path PE3-PE1
to 42.42.42.42
!
rsvp-trunk PE3-PE2 ipv4
to 30.30.30.30
!
rsvp-trunk PE3-PE4 ipv4
to 19.19.19.19
!
!
end

!
```

```
PE3-7028#sh ip ospf neighbor
```

```

Total number of full neighbors: 2
OSPF process 100 VRF(default):
Neighbor ID      Pri   State           Dead Time   Address      Interface     Instance ID
11.11.11.11      1    Full/Backup     00:00:37    30.30.40.1   xe21          0
9.9.9.9          1    Full/Backup     00:00:29    40.40.40.1   xe11          0
PE3-7028#
PE3-7028#sh rsvp session
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
* indicates the session is active with local repair at one or more nodes
(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime    Rt   Style  Labelin  Labelout
42.42.42.42  28.28.28.28  5001     2201     PRI     PE3-PE1-         24330
Primary          UP    00:22:34  1 1 SE    -
30.30.30.30  28.28.28.28  5002     2202     PRI     PE3-PE2-         24320
Primary          UP    00:27:28  1 1 SE    -
19.19.19.19  28.28.28.28  5003     2203     PRI     PE3-PE4-         24324
Primary          UP    00:24:33  1 1 SE    -
Total 3 displayed, Up 3, Down 0.

Egress RSVP:
To          From          Tun-ID    LSP-ID    Type    LSPName          State
Uptime    Rt   Style  Labelin  Labelout
28.28.28.28  30.30.30.30  5002     2201     PRI     PE2-PE3-         24960
Primary          UP    00:27:22  1 1 SE    -
28.28.28.28  42.42.42.42  5002     2203     PRI     PE1-PE3-         24962
Primary          UP    00:24:01  1 1 SE    -
28.28.28.28  19.19.19.19  5003     2203     PRI     PE4-PE3-         24961
Primary          UP    00:24:57  1 1 SE    -
Total 3 displayed, Up 3, Down 0.

PE3-7028#
PE3-7028#
PE3-7028#show bgp l2vpn evpn
BGP table version is 8, local router ID is 28.28.28.28
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

Network      Next Hop      Metric    LocPrf      Weight    Path  Peer      Encap
RD[65100:501]
*>i   [3]:[501]:[32,30.30.30.30]
              30.30.30.30      0        100        0    i      -  11.11.11.11    MPLS
RD[65100:502]
*>i   [3]:[502]:[32,30.30.30.30]
              30.30.30.30      0        100        0    i      -  11.11.11.11    MPLS
RD[65100:503]
*>i   [3]:[503]:[32,30.30.30.30]
              30.30.30.30      0        100        0    i      -  11.11.11.11    MPLS
RD[65100:504]
*>i   [3]:[504]:[32,30.30.30.30]
              30.30.30.30      0        100        0    i      -  11.11.11.11    MPLS

```

```

RD[65100:505]
*>i    [3]:[505]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:506]
*>i    [3]:[506]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:507]
*>i    [3]:[507]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:508]
*>i    [3]:[508]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:509]
*>i    [3]:[509]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:510]
*>i    [3]:[510]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:250]
*>i    [5]:[0]:[0]:[24]:[70.70.1.0]:[0.0.0.0]:[16]
          19.19.19.19          0          100          0    ?          -  11.11.11.11          MPLS

RD[19.19.19.19:801]
*>i    [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
          19.19.19.19          0          100          0    ?          -  11.11.11.11          MPLS

RD[19.19.19.19:3000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [3]:[251]:[32,19.19.19.19]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:4000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [3]:[811]:[32,19.19.19.19]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:64512]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS
*>i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
          19.19.19.19          0          100          0    i          -  11.11.11.11          MPLS

RD[28.28.28.28:2000] VRF[vrfirb]
*>    [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
          28.28.28.28          0          100          32768    i          -  -----
MPLS
* i    [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]

```

[illegible]

[illegible]


```

42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0047]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0048]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0049]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:004a]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:004b]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:004c]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:004d]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:004e]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:004f]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0050]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0051]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0052]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0053]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
* i      [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>      [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
28.28.28.28      0      100      32768      i      - -----
MPLS
*>      [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
28.28.28.28      0      100      32768      i      - -----
MPLS
* i      [3]:[200]:[32,42.42.42.42]

```

```

42.42.42.42          0          100      0    i      - 11.11.11.11    MPLS
*> [3]:[201]:[32,28.28.28.28]
      28.28.28.28      0          100    32768  i      - -----
MPLS

RD[28.28.28.28:4000] VRF[vrfirbMH801]
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      19.19.19.19      0          100      0    i      - 11.11.11.11    MPLS
*>      28.28.28.28      0          100    32768  i      - -----
MPLS
* i [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      19.19.19.19      0          100      0    i      - 11.11.11.11    MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
* i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
* i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      19.19.19.19      0          100      0    i      - 11.11.11.11    MPLS
*>      28.28.28.28      0          100    32768  i      - -----
MPLS
* i [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      19.19.19.19      0          100      0    i      - 11.11.11.11    MPLS
*>      28.28.28.28      0          100    32768  i      - -----
MPLS
* i [3]:[801]:[32,42.42.42.42]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
* i [3]:[811]:[32,19.19.19.19]
      19.19.19.19      0          100      0    i      - 11.11.11.11    MPLS
*> [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0          100    32768  i      - -----
MPLS

RD[28.28.28.28:64512] VRF[evpn-gvrf-1]
*> [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0          100    32768  i      - -----
MPLS
* i [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
      19.19.19.19      0          100      0    i      - 11.11.11.11    MPLS
*> [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0          100    32768  i      - -----
MPLS

RD[30.30.30.30:250]
*>i [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0          100      0    ?      - 11.11.11.11    MPLS

RD[30.30.30.30:3000]
*>i [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0          100      0    i      - 11.11.11.11    MPLS
*>i [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
      30.30.30.30      0          100      0    i      - 11.11.11.11    MPLS
*>i [3]:[250]:[32,30.30.30.30]
      30.30.30.30      0          100      0    i      - 11.11.11.11    MPLS

RD[42.42.42.42:801]
*>i [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
      42.42.42.42      0          100      0    ?      - 11.11.11.11    MPLS

RD[42.42.42.42:2000]
*>i [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
      42.42.42.42      0          100      0    i      - 11.11.11.11    MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]

```

[illegible]

[illegible]

```

42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0047]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0048]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0049]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004a]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004b]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004c]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004d]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004e]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004f]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0050]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0051]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0052]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0053]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [3]:[200]:[32,42.42.42]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS

RD[42.42.42.42:4000]
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
42.42.42.42      0      100      0      i      -      11.11.11.11      MPLS
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]

```

```

      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i  [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i  [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i  [3]:[801]:[32,42.42.42.42]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS

```

Total number of prefixes 254

PE3-7028#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN

(m) - Service mapped over multipath transport

(e) - Service mapped over ECMP

(D) - Down

Ext-Color - Extended-community color advertised by BGP

B(x) - BGP EVPN MPLS Services

Code	FEC	Out-Intf	Nexthop	UpTime	FTN-ID	VRF-ID	Nhlfe-ID	Pri	Out-
Label	Out-Intf	Nexthop	UpTime	Ext-Color					
B(x)>	81.81.1.0/24				1	3	7	-	-
	-		00:22:37	-			5	Yes	17
	42.42.42.42	-	-						-

PE3-7028#show ip route vrf all database bgp

IP Route Table for VRF "default"

Total number of IPv4 routes 0

IP Route Table for VRF "management"

Total number of IPv4 routes 0

IP Route Table for VRF "ip_vrfirb"

Total number of IPv4 routes 0

IP Route Table for VRF "ip_vrfirbMH801"

B > 81.81.1.0/24 [200/0] via 42.42.42.42, installed 00:23:07, last update 00:23:07 ago

Total number of IPv4 routes 1

IP Route Table for VRF "vrfirb"

Total number of IPv4 routes 0

IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0

IP Route Table for VRF "evpn-gvrf-1"

Total number of IPv4 routes 0

Gateway of last resort is not set

PE3-7028#show hsl mpls l3vpn-fts

TABLE: L3VPN Ftn table

Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup

L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy

VRF	DESTINATION	TNL/LSP/TYPE/NHLFE-IX	EGRESS	UpTime	Ref	Prefix
MPLS	OUT	OUT	NEXTHOP			
ID	FEC	Ifname	IFNAME	LABEL		
OBJECT	cnt	cnt				
3	B> 81.81.1.0/24	7(s)	-	-	-	-
	0x2000ccdc	00:22:37	1	1		
	-	/-	/PRI	/5		
	17	42.42.42.42	0x2000ccdc			

PE3-7028#sh evpn mpls tunnel

EVPN-MPLS Network tunnel Entries

Source	Destination	Status	Up/Down	Update	evpn-id	Local-
Leaf	Remote-Leaf	Ext-Color	FAT			

```
=====
=====
28.28.28.28      42.42.42.42      Installed      00:23:09      00:23:09      40000      ---
    ---          ---          ---
28.28.28.28      19.19.19.19      Installed      00:24:38      00:24:38      811        ---
    ---          ---          ---
28.28.28.28      19.19.19.19      Installed      00:25:03      00:25:03      40000      ---
    ---          ---          ---

Total number of entries are 3
PE3-7028#
apps-fileview.texmex_20251030.00_p0
PE3.txt
Displaying PE3.txt.
```

PE4

```
PE4-7019#sh running-config
!
! Software version: EC_AS5916-54X-OcNOS-SP-MPLS-7.0.0.136-
Alpha                                                    10/12/2025 17:36:53
!
! Last configuration change at 12:42:34 UTC Wed Oct 15 2025 by
r                                                         oot
!
!
service password-encryption
!
logging console 5
logging logfile rsvp 7
logging level nsm 4
logging level rsvp 7
logging level bgp 7
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile filter evpn-mpls-mh enable
hardware-profile statistics ingress-acl enable
!
qos enable
!
hostname PE4-7019
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn mpls irb
!
evpn mpls multihoming enable
```

```
!  
ip vrf ip_vrfirb250  
  rd 19.19.19.19:250  
  route-target both 250:250  
  l3vni 30000  
!  
ip vrf ip_vrfirbMH801  
  rd 19.19.19.19:801  
  route-target both 801:801  
  l3vni 40000  
!  
ip vrf management  
!  
mac vrf vrfirb250  
  rd 19.19.19.19:3000  
  route-target both 3000:3000  
!  
mac vrf vrfirbMH801  
  rd 19.19.19.19:4000  
  route-target both 4000:4000  
!  
evpn mpls vtep-ip-global 19.19.19.19  
!  
evpn mpls id 251  
  host-reachability-protocol evpn-bgp vrfirb250  
  evpn irb irb151  
!  
evpn mpls id 811  
  host-reachability-protocol evpn-bgp vrfirbMH801  
  evpn irb irb711  
!  
router rsvp  
  auto-bypass  
  attributes best-effort  
  protection-capability node  
  reoptimize  
  exit  
  inactivity-timer 5  
  enable  
  exit  
!  
interface po20  
  switchport  
  evpn multi-homed system-mac 0000.1111.2222  
!  
interface po20.251 switchport  
  encapsulation dot1q 251  
  rewrite pop  
  access-if-evpn  
  map vpn-id 251  
!  
interface po20.811 switchport  
  encapsulation dot1q 811  
  rewrite pop  
  access-if-evpn  
  map vpn-id 811  
!  
interface ce0  
!  
interface ce1  
!  
interface ce2  
!  
interface ce3  
!  
interface ce4  
!  
interface ce5
```



```
!  
interface eth0  
    ip vrf forwarding management  
    ip address dhcp  
!  
interface irb151  
    ip vrf forwarding ip_vrfirb250  
    ip address 70.70.1.1/24 anycast  
    ipv6 address 70::1/64 anycast  
!  
interface irb711  
    ip vrf forwarding ip_vrfirbMH801  
    evpn irb-if-forwarding anycast-gateway-mac  
    ip address 91.91.1.1/24 anycast  
    ipv6 address 91::1/64 anycast  
    ip ospf network point-to-point  
!  
interface lo  
    ip address 127.0.0.1/8  
    ip address 19.19.19.19/32 secondary  
    ipv6 address ::1/128  
    ip router isis 1  
!  
interface lo.management  
    ip vrf forwarding management  
    ip address 127.0.0.1/8  
    ipv6 address ::1/128  
!  
interface xe0  
!  
interface xe1  
!  
interface xe2  
!  
interface xe3  
    channel-group 20 mode active  
!  
interface xe4  
!  
interface xe5  
!  
interface xe6  
!  
interface xe7  
!  
interface xe8  
!  
interface xe9  
!  
interface xe10  
!  
interface xe11  
!  
interface xe12  
!  
interface xe13  
!  
interface xe14  
!  
interface xe15  
!  
interface xe16  
!  
interface xe17  
!  
interface xe18  
!  
interface xe19
```

```
!  
interface xe20  
!  
interface xe21  
!  
interface xe22  
!  
interface xe23  
!  
interface xe24  
!  
interface xe25  
!  
interface xe26  
!  
interface xe27  
!  
interface xe28  
!  
interface xe29  
!  
interface xe30  
!  
interface xe31  
!  
interface xe32  
!  
interface xe33  
!  
interface xe34  
!  
interface xe35  
!  
interface xe36  
!  
interface xe37  
!  
interface xe38  
!  
interface xe39  
!  
interface xe40  
!  
interface xe41  
!  
interface xe42  
!  
interface xe43  
!  
interface xe44  
  load-interval 30  
  ip address 30.30.30.2/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-rsvp  
!  
interface xe45  
!  
interface xe46  
  load-interval 30  
  ip address 40.40.50.2/24  
  mtu 9216  
  label-switching  
  ip router isis 1  
  enable-rsvp  
!  
interface xe47
```

```

!
evpn irb-forwarding anycast-gateway-mac 0011.2200.0002
!
exit
!
router ospf 100
bfd all-interfaces
network 19.19.19.19/32 area 0.0.0.0
network 30.30.30.0/24 area 0.0.0.0
network 40.40.50.0/24 area 0.0.0.0
!
router bgp 65100
bgp router-id 19.19.19.19
no bgp inbound-route-filter
neighbor 11.11.11.11 remote-as 65100
neighbor 11.11.11.11 update-source lo
neighbor 91.81.1.2 remote-as 65102
!
address-family l2vpn evpn
neighbor 11.11.11.11 activate
neighbor 11.11.11.11 route-reflector-client
exit-address-family
!
address-family ipv4 vrf ip_vrfirb250
redistribute connected
exit-address-family
!
address-family ipv4 vrf ip_vrfirbMH801
max-paths ibgp 4
redistribute connected
redistribute static
neighbor 91.91.1.2 remote-as 65102
neighbor 91.91.1.2 activate
exit-address-family
!
exit
!
rsvp-path PE4-PE1 mpls
30.30.30.1 strict
10.10.10.1 strict
!
rsvp-trunk PE4-PE1 ipv4
no primary cspf
primary path PE4-PE1
to 42.42.42.42
!
rsvp-trunk PE4-PE2 ipv4
to 30.30.30.30
!
rsvp-trunk PE4-PE3 ipv4
to 28.28.28.28
!

```

```
PE4-7019#sh ip ospf neighbor
```

```
Total number of full neighbors: 2
```

```
OSPF process 100 VRF(default):
```

Neighbor ID	Pri	State	Dead Time	Address	Interface	Instance ID
11.11.11.11	1	Full/Backup	00:00:37	30.30.30.1	xe44	0
9.9.9.9	1	Full/Backup	00:00:36	40.40.50.1	xe46	0

```
PE4-7019#
```

```
PE4-7019#sh rsvp session
```

```
Type : PRI - Primary, SEC - Secondary, DTR - Detour, BPS - Bypass
```

```
State : UP - Up, DN - Down, BU - Backup in Use, SU - Secondary in Use, FS - Forced to Secondary
```

```
* indicates the session is active with local repair at one or more nodes
```

(P) indicates the secondary-priority session is acting as primary

Ingress RSVP:

To	From	Tun-ID	LSP-ID	Type	LSPName	State
Uptime	Rt	Style	Labelin	Labelout		
42.42.42.42	19.19.19.19	5001	2201	PRI	PE4-PE1-	
Primary	UP	00:23:52	1 1 SE	-	24328	
30.30.30.30	19.19.19.19	5002	2202	PRI	PE4-PE2-	
Primary	UP	00:25:41	1 1 SE	-	24323	
28.28.28.28	19.19.19.19	5003	2203	PRI	PE4-PE3-	
Primary	UP	00:25:50	1 1 SE	-	24322	

Total 3 displayed, Up 3, Down 0.

Egress RSVP:

To	From	Tun-ID	LSP-ID	Type	LSPName	State
Uptime	Rt	Style	Labelin	Labelout		
19.19.19.19	28.28.28.28	5003	2203	PRI	PE3-PE4-	
Primary	UP	00:25:26	1 1 SE	24960	-	
19.19.19.19	42.42.42.42	5003	2204	PRI	PE1-PE4-	
Primary	UP	00:24:54	1 1 SE	24961	-	

Total 2 displayed, Up 2, Down 0.

PE4-7019#

PE4-7019#

PE4-7019#show bgp l2vpn evpn

BGP table version is 6, local router ID is 19.19.19.19

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal,

l - labeled, S Stale

Origin codes: i - IGP, e - EGP, ? - incomplete

Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]

1 - Ethernet Auto-discovery Route

2 - MAC/IP Route

3 - Inclusive Multicast Route

4 - Ethernet Segment Route

5 - Prefix Route

Network	Next Hop	Metric	LocPrf	Weight	Path	Peer	Encap
RD[65100:501]							
*>i [3]:[501]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:502]							
*>i [3]:[502]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:503]							
*>i [3]:[503]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:504]							
*>i [3]:[504]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:505]							
*>i [3]:[505]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:506]							
*>i [3]:[506]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS
RD[65100:507]							
*>i [3]:[507]:[32,30.30.30.30]	30.30.30.30	0	100	0 i	-	11.11.11.11	MPLS

```

RD[65100:508]
*>i    [3]:[508]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:509]
*>i    [3]:[509]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:510]
*>i    [3]:[510]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:3000] VRF[vrfirb250]
*>    [1]:[00:00:00:11:11:22:22:00:00]:[251]:[19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS
*>    [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
*>    [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1]:[19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [3]:[250]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS
*>    [3]:[251]:[32,19.19.19.19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS

RD[19.19.19.19:4000] VRF[vrfirbMH801]
* i    [1]:[00:00:00:11:11:22:22:00:00]:[811]:[20]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
*>    [1]:[00:00:00:11:11:22:22:00:00]:[811]:[20]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [1]:[00:00:00:11:11:22:22:00:00]:[4294967295]:[0]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
*>    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
*>    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [3]:[801]:[32,42.42.42.42]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
*>    [3]:[811]:[32,19.19.19.19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [3]:[811]:[32,28.28.28.28]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:64512] VRF[evpn-gvrf-1]
*>    [1]:[00:00:00:11:11:22:22:00:00]:[4294967295]:[0]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
*>    [4]:[00:00:00:11:11:22:22:00:00]:[32,19.19.19.19]
          19.19.19.19          0          100          32768    i          -  -----

```

```

MPLS
* i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS

RD[28.28.28.28:200]
*>i    [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
      28.28.28.28      0      100      0      ?      - 11.11.11.11      MPLS

RD[28.28.28.28:801]
*>i    [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
      28.28.28.28      0      100      0      ?      - 11.11.11.11      MPLS

RD[28.28.28.28:2000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [3]:[201]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS

RD[28.28.28.28:4000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS

RD[28.28.28.28:64512]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS
*>i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      - 11.11.11.11      MPLS

RD[30.30.30.30:250]
*>i    [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0      100      0      ?      - 11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS
*>i    [3]:[250]:[32,30.30.30.30]
      30.30.30.30      0      100      0      i      - 11.11.11.11      MPLS

RD[42.42.42.42:801]
*>i    [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
      42.42.42.42      0      100      0      ?      - 11.11.11.11      MPLS

RD[42.42.42.42:2000]
*>i    [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
      42.42.42.42      0      100      0      i      - 11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]

```

[illegible]

[illegible]


```

42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004b]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004c]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004d]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004e]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004f]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0050]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0051]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0052]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0053]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[200]:[32,42.42.42.42]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS

RD[42.42.42.42:4000]
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1]:[19]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1]:[19]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[801]:[32,42.42.42.42]
42.42.42.42 0 100 0 i - 11.11.11.11 MPLS

```

```

Total number of prefixes 155
PE4-7019#show mpls vrf-forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN
(m) - Service mapped over multipath transport
(e) - Service mapped over ECMP
(D) - Down
Ext-Color - Extended-community color advertised by BGP
B(x) - BGP EVPN MPLS Services

Code    FEC          Nexthop      UpTime      FTN-ID VRF-ID    Nhlfe-ID    Pri    Out-
Label    Out-Intf          Nexthop      UpTime      Ext-Color
B(x)>    60.60.1.0/24
        -              00:25:57 -
        -              3
e44      (D) 30.30.30.30    -          -
B(x)>    81.81.1.0/24
        -              00:23:54 -
        -              1
        42.42.42.42    -          -
        -              3
        7
        Yes 16
        Yes 17
        -
PE4-7019#show ip route vrf all database bgp
IP Route Table for VRF "default"

Total number of IPv4 routes 0
IP Route Table for VRF "management"

Total number of IPv4 routes 0
IP Route Table for VRF "ip_vrfirb250"
B      > 60.60.1.0/24 [200/0] via 30.30.30.30, installed 00:25:57, last update 00:25:57 ago

Total number of IPv4 routes 1
IP Route Table for VRF "ip_vrfirbMH801"
B      > 81.81.1.0/24 [200/0] via 42.42.42.42, installed 00:24:00, last update 00:24:00 ago

Total number of IPv4 routes 1
IP Route Table for VRF "vrfirb250"

Total number of IPv4 routes 0
IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0
IP Route Table for VRF "evpn-gvrf-1"

Total number of IPv4 routes 0

Gateway of last resort is not set
PE4-7019#show hsl mpls l3vpn-ftn
TABLE: L3VPN Ftn table
Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup
L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| VRF | DESTINATION | TNL/LSP/TYPE/NHLFE-IX | EGRESS | UpTime | Ref | Prefix |
| MPLS | OUT | OUT | NEXTHOP | | | |
| ID | FEC | cnt | cnt | Ifname | IFNAME | LABEL |
| OBJECT | | | | | | |
+-----+-----+-----+-----+-----+-----+-----+-----+
3    B> 81.81.1.0/24
      0x2000101a 00:23:54 1 1 9(s) - - -
      - /- /PRI /7 -
      17 42.42.42.42 0x2000101a
PE4-7019#sh evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color FAT
=====
19.19.19.19 28.28.28.28      Installed   00:25:54     00:25:54     811          ---
  ---          ---

```

```

19.19.19.19      28.28.28.28      Installed      00:25:59      00:25:59      40000      ---
---
19.19.19.19      42.42.42.42      Installed      00:24:02      00:24:02      40000      ---
---
19.19.19.19      30.30.30.30      Installed      00:25:59      00:25:59      30000      ---
---

Total number of entries are 4
PE4-7019#
apps-fileview.texmex_20251030.00_p0
PE4.txt
Displaying PE4.txt.

```

Validation

To verify the ECMP configuration, check the output of the following `show` commands:

```

PE4-7019#show bgp l2vpn evpn
BGP table version is 6, local router ID is 19.19.19.19
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route

      Network          Next Hop          Metric    LocPrf      Weight    Path  Peer          Encap
RD[65100:501]
*>i    [3]:[501]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:502]
*>i    [3]:[502]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:503]
*>i    [3]:[503]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:504]
*>i    [3]:[504]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:505]
*>i    [3]:[505]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:506]
*>i    [3]:[506]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:507]
*>i    [3]:[507]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS
RD[65100:508]
*>i    [3]:[508]:[32,30.30.30.30]
              30.30.30.30          0          100         0    i          -  11.11.11.11    MPLS

```

```

RD[65100:509]
*>i    [3]:[509]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[65100:510]
*>i    [3]:[510]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:3000] VRF[vrfirb250]
*>    [1]:[00:00:00:11:11:22:22:00:00:00]:[251]:[19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1][19]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS
*>    [2]:[0]:[251]:[48,d077:ceaa:8001]:[32,70.70.1.1]:[19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
*>    [2]:[0]:[251]:[48,d077:ceaa:8001]:[128,70::1][19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [3]:[250]:[32,30.30.30.30]
          30.30.30.30          0          100          0    i          -  11.11.11.11          MPLS
*>    [3]:[251]:[32,19.19.19.19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS

RD[19.19.19.19:4000] VRF[vrfirbMH801]
* i    [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
*>    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
          28.28.28.28          0          100          32768    i          -  -----
MPLS
* i    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1][19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1][19]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
* i    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS
*>    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1][20]
          28.28.28.28          0          100          32768    i          -  -----
MPLS
* i    [3]:[801]:[32,42.42.42.42]
          42.42.42.42          0          100          0    i          -  11.11.11.11          MPLS
*>    [3]:[811]:[32,19.19.19.19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [3]:[811]:[32,28.28.28.28]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS

RD[19.19.19.19:64512] VRF[evpn-gvrf-1]
*>    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
*>    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,19.19.19.19]
          19.19.19.19          0          100          32768    i          -  -----
MPLS
* i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
          28.28.28.28          0          100          0    i          -  11.11.11.11          MPLS

```

```

RD[28.28.28.28:200]
*>i    [5]:[0]:[0]:[24]:[90.90.1.0]:[0.0.0.0]:[16]
      28.28.28.28      0      100      0      ?      -  11.11.11.11      MPLS

RD[28.28.28.28:801]
*>i    [5]:[0]:[0]:[24]:[91.91.1.0]:[0.0.0.0]:[17]
      28.28.28.28      0      100      0      ?      -  11.11.11.11      MPLS

RD[28.28.28.28:2000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[201]:[19]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[201]:[48,0090:fb7d:ad12]:[32,90.90.1.1]:[19]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[201]:[48,0090:fb7d:ad12]:[128,90::1]:[19]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [3]:[201]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS

RD[28.28.28.28:4000]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[811]:[20]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[32,91.91.1.1]:[20]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[811]:[48,0011:2200:0002]:[128,91::1]:[20]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [3]:[811]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS

RD[28.28.28.28:64512]
*>i    [1]:[00:00:00:11:11:22:22:00:00:00]:[4294967295]:[0]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS
*>i    [4]:[00:00:00:11:11:22:22:00:00:00]:[32,28.28.28.28]
      28.28.28.28      0      100      0      i      -  11.11.11.11      MPLS

RD[30.30.30.30:250]
*>i    [5]:[0]:[0]:[24]:[60.60.1.0]:[0.0.0.0]:[16]
      30.30.30.30      0      100      0      ?      -  11.11.11.11      MPLS

RD[30.30.30.30:3000]
*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[32,60.60.1.1]:[19]
      30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[250]:[48,b86a:97c6:33c5]:[128,60::1]:[19]
      30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS
*>i    [3]:[250]:[32,30.30.30.30]
      30.30.30.30      0      100      0      i      -  11.11.11.11      MPLS

RD[42.42.42.42:801]
*>i    [5]:[0]:[0]:[24]:[81.81.1.0]:[0.0.0.0]:[17]
      42.42.42.42      0      100      0      ?      -  11.11.11.11      MPLS

RD[42.42.42.42:2000]
*>i    [2]:[0]:[200]:[48,1200:0000:0000]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0001]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0002]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0003]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0004]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0005]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0006]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS
*>i    [2]:[0]:[200]:[48,1200:0100:0007]:[0]:[18]
      42.42.42.42      0      100      0      i      -  11.11.11.11      MPLS

```

© 2026 IP Infusion Inc. Proprietary 94

© 2026 IP Infusion Inc. Proprietary 95

```

*>i [2]:[0]:[200]:[48,1200:0100:004c]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004d]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004e]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:004f]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0050]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0051]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0052]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0053]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0054]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0055]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0056]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0057]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0058]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0059]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005a]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005b]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005c]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005d]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005e]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:005f]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0060]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0061]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0062]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0063]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[200]:[48,1200:0100:0064]:[0]:[18]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[200]:[32,42.42.42]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS

RD[42.42.42.42:4000]
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[0]:[19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,0010:9400:0003]:[128,fe80::1][19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[32,81.81.1.1]:[19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [2]:[0]:[801]:[48,5c07:5854:1200]:[128,81::1][19]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS
*>i [3]:[801]:[32,42.42.42]
    42.42.42.42 0 100 0 i - 11.11.11.11 MPLS

```

Total number of prefixes 155

PE4-7019#show mpls vrf-forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup, B - BGP FTN

(m) - Service mapped over multipath transport

(e) - Service mapped over ECMP

(D) - Down

Ext-Color - Extended-community color advertised by BGP

B(x) - BGP EVPN MPLS Services

Code	FEC	Out-Intf	Nexthop	UpTime	FTN-ID	VRF-ID	Nhlfe-ID	Pri	Out-
Label					Ext-Color				
B(x)>	60.60.1.0/24				1	2	4	-	-
	-		00:25:57	-			3	Yes	16
e44	(D) 30.30.30.30		-	-					x
B(x)>	81.81.1.0/24				1	3	9	-	-
	-		00:23:54	-			7	Yes	17
	42.42.42.42		-	-					-

PE4-7019#show ip route vrf all database bgp

IP Route Table for VRF "default"

Total number of IPv4 routes 0

IP Route Table for VRF "management"

Total number of IPv4 routes 0

IP Route Table for VRF "ip_vrfirb250"

B > 60.60.1.0/24 [200/0] via 30.30.30.30, installed 00:25:57, last update 00:25:57 ago

Total number of IPv4 routes 1

IP Route Table for VRF "ip_vrfirbMH801"

B > 81.81.1.0/24 [200/0] via 42.42.42.42, installed 00:24:00, last update 00:24:00 ago

Total number of IPv4 routes 1

IP Route Table for VRF "vrfirb250"

Total number of IPv4 routes 0

IP Route Table for VRF "vrfirbMH801"

Total number of IPv4 routes 0

IP Route Table for VRF "evpn-gvrf-1"

Total number of IPv4 routes 0

Gateway of last resort is not set

PE4-7019#show hsl mpls l3vpn-fts

TABLE: L3VPN Ftn table

Codes: > - installed FTN, (e) - ecmp, (s) - single(non-ecmp), (p) - primary, (b) - backup

L - LDP, K - Static, R - RSVP, B - BGP, O - OSPF-SR, I - ISIS-SR, P - SR-Policy

VRF	DESTINATION	TNL/LSP/TYPE/NHLFE-IX	EGRESS	UpTime	Ref	Prefix
MPLS	OUT	OUT	NEXTHOP			
ID	FEC		Ifname	IFNAME	LABEL	
OBJECT		cnt	cnt			
3	B> 81.81.1.0/24	9(s)	-	-	-	-
	0x2000101a	00:23:54	1	1		
	-	/-	/PRI	/7	-	
	17	42.42.42.42	0x2000101a			

PE4-7019#sh evpn mpls tunnel

EVPN-MPLS Network tunnel Entries

Source	Destination	Status	Up/Down	Update	evpn-id	Local-
Leaf	Remote-Leaf	Ext-Color	FAT			

19.19.19.19	28.28.28.28	Installed	00:25:54	00:25:54	811	---
19.19.19.19	28.28.28.28	Installed	00:25:59	00:25:59	40000	---
19.19.19.19	42.42.42.42	Installed	00:24:02	00:24:02	40000	---

```

      ---      ---
19.19.19.19    30.30.30.30    Installed    00:25:59    00:25:59    30000    ---
      ---      ---

Total number of entries are 4
PE4-7019#

```

ECMP for L3EVPN Commands

The ECMP feature in L3EVPN has the following configuration commands:



Note: The `multipath` CLI is available under the EVPN address family. When this command is configured, BGP computes multiple best paths per Route Distinguisher (RD) and advertises these routes to the associated VRFs. The VRFs then install the routes based on the `max-paths` configuration defined at the VRF level. If the `multipath` CLI is not configured and only `max-paths` is configured under the VRF, BGP propagates only the single best route from the RD to the VRF. The VRF then installs multiple paths based on the `max-paths` configuration.

multi-path

Use this command to configure BGP to calculate the multiple routes at RD level for a given address family. Currently, OcNOS supports this command for `vpn4` and `vpn6` unicast address families. When this command is configured, BGP processes the multiple best routes per RD and publishes the routes to multiple VRFs, where the routes get installed based on the BGP config `<max-paths>` at VRF level.

It is recommended to use this command along with `config <max-paths>` at VRF level. This facilitates the ECMP functionality for VPN routes so the best path calculations are done at each RD level and VRF level and the final ECMP paths are chosen at VRF level.

Command Syntax

```

multi-path
no multi-path

```

Parameters

None

Default

Disabled

Command Mode

BGP Address-family L2VPN EVPN

Examples

Consider a topology such that BGP on local PE receives four advertisements from each of its peer PEs, with below IP-VRF parameters:

-RT=200:1, RD=x from PE-1

-RT=200:1, RD=x from PE-2

-RT=200:1, RD=y from PE-3

-RT=200:1, RD=y from PE-4.

Example 1

The following example illustrates the installation of four paths in VRF RED.

Configure a VRF

```
7535-2(config)#ip vrf RED
7535-2(config-vrf)#rd 2:2
7535-2(config-vrf)#route-target both 200:1
7535-2(config-vrf)#exit
```

Configure BGP

```
7535-2(config)#router bgp 100
7535-2(config-router)#address-family l2vpn evpn
7535-2(config-router-af)#multi-path
7535-2(config-router-af)#exit
7535-2(config-router)#address-family ipv4 vrf RED
7535-2(config-router-af)#max-paths ibgp 4
```

Here, BGP learns 2 ECMP paths with RT 200:1 with RD x, and 2 ECMP paths with RT 200:1 with RD y. Finally, VRF RED will have a route with four multipaths installed in the fastpath and this leads to load balancing of user traffic across the four paths.

Example 2

The below example installs two paths in VRF RED. The Multipath configuration is not enabled, however `<max-paths>` is configured at VRF level.

```
7535-2(config)#router bgp 100
7535-2(config-router)#address-family l2vpn evpn
7535-2(config-router-af)#exit
7535-2(config-router)#address-family ipv4 vrf RED
7535-2(config-router-af)#max-paths ibgp 4
```

Here, BGP learns 2 ECMP paths with RT 200:1 with RD x, and 2 ECMP paths with RT 200:1 with RD y. BGP will not calculate the multipaths per RD, and it will select only the best (one path) per RD. Finally, VRF RED will have a route with two multipaths installed in the fastpath and this leads to load balancing of user traffic across the two paths.

To unconfigure the BGP EVPN Multipath functionality, both `<multi-path>` and `<max-path>` configurations must be removed.

EVPN AF Route-Maps and Route Filtering

Overview

The EVPN AF Route-Maps and Route Filtering feature extends the existing route-map framework to the L2VPN EVPN address family in BGP. It allows users to filter, modify, or manage EVPN routes exchanged between BGP peers.

A route-map defines the criteria to match specific route attributes and actions to set or modify those attributes. Route-maps can be applied in the incoming (IN) or outgoing (OUT) direction on a BGP neighbor or peer group.

- OUT direction: Processes EVPN routes before BGP advertises them.
- IN direction: Processes routes after BGP receives them.

The framework introduces EVPN-specific match conditions, such as route type and MAC lists, enabling precise control over route propagation and reducing control-plane overhead in large BGP-EVPN networks.

Feature Characteristics

- Supports route-map configuration under the L2VPN EVPN address family.
- Applies route-maps in both IN and OUT directions.
- Integrates with the existing route-map framework with EVPN extensions.
- Supports numbered and unnumbered L2VPN modes.
- Evaluates routes using match and set conditions.
- Adds EVPN-specific match capabilities:
 - match evpn-route-type : Type-1 to Type-5, MAC-only, MAC-IP routes.
- Match mac-list : Permit or deny specific MAC addresses.
- Supports standard BGP route-map matches like ASpath, next-hop, route-target and other match options.
- Controls route propagation direction and scope.
- Reduces unnecessary route advertisements across spines, leaves, and superspines.

Benefits

- Provides granular EVPN route control.
- Allows selective advertisement, acceptance, or rejection of routes.
- Enhances scalability and performance by reducing redundant updates.
- Simplifies multi-tenant isolation via communities or route-targets.
- Ensures consistent configuration across IPv4, IPv6, and EVPN AFs.
- Reduces control-plane load and processing time for operational efficiency.

Prerequisites

Before configuring EVPN AF Route-Maps and Route Filtering, ensure the following conditions are met:

- BGP EVPN Setup:
 - BGP sessions between all relevant PEs, RRs, and leaf/spine nodes must be established.
 - The L2VPN EVPN address family must be enabled on all participating BGP neighbors.

Limitation

- The set option is not supported for import-map in VRF configurations.
- The AIGP path attribute is not supported for EVPN address-family routes. Therefore, when configuring a neighbor route map under the EVPN address family or a route map used in a VRF export map, the **set aigp-metric <>** command is not applicable to EVPN routes.
- Under route-map configuration, **set ip nexthop** is not supported for EVPN address family routes.

Configuration

This section describes the configuration of EVPN Address Family (AF) with route filtering using route-maps on PE routers.

Route-maps are applied in the EVPN AF to control route import or export policies and to modify all the supported attributes.

Topology

The topology illustrates an EVPN network using MPLS underlay and BGP control plane for end-to-end connectivity between CE1 and CE2.

Figure 3. EVPN AF Route Map



PE1 and PE3 function as Provider Edge routers participating in EVPN BGP sessions, while RR acts as the Route Reflector to distribute EVPN routes between them.

PE1–RR uses LDP with OSPF, and RR–PE3 uses RSVP with ISIS for MPLS transport. BGP Labeled Unicast (BGP-LU) establishes end-to-end label distribution between PEs.

Route-maps are applied for route control:

- PE1 applies an OUT route-map for route advertisement.
- RR applies an IN route-map for route reflection.
- PE3 applies an IN route-map for route import and CE2 advertisement.

Traffic from CE1 to CE2 traverses the MPLS core, with BGP EVPN managing MAC and IP route exchange between PEs.

Router or Node Configuration Steps

Perform the following configuration steps to apply route-maps under the BGP EVPN address family, enabling selective advertisement and acceptance of EVPN routes between peers.



Note: Before configuration meet all [Prerequisites \(page 100\)](#).

1. Configure MAC Access List on PE1.

Define a MAC access list to restrict specific MAC addresses from being advertised through EVPN.

```
mac-list mac_list1
seq 10 deny 6821.5f1f.5220 ffff.ffff.ffff
seq 20 permit 0000.0000.0000 0000.0000.0000
```

2. Configure Outbound Route-Map on PE1.

a. Define a route-map named **rm_evpn_out** to control the EVPN route advertisement.

b. Match different EVPN route types and set appropriate BGP attributes for each type.

```
route-map rm_evpn_out permit 10
match evpn-route-type type-1
set local-preference 200
set community 65000:100 additive

route-map rm_evpn_out permit 20
match evpn-route-type type-2
match mac address list mac_list1
set local-preference 150
set metric 50
set community 65000:200 additive

route-map rm_evpn_out permit 30
match evpn-route-type type-3
set local-preference 100
set origin igp
set community 65000:300 additive

route-map rm_evpn_out permit 1000
```

c. Apply the route-map to the BGP EVPN address family for outbound updates to the Route Reflector.

```
router bgp 65000
address-family l2vpn evpn
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 route-map rm_evpn_out out
exit-address-family
```

3. Configure Inbound Route-Map on PE3.

a. Create a route-map named **rm_evpn_in** to control inbound EVPN route processing.

b. Match specific EVPN route types and adjust route selection parameters.

```
route-map rm_evpn_in permit 10
match evpn-route-type type-1
set weight 200
```

c. Apply the route-map to the BGP EVPN address family for inbound routes from the Route Reflector.

```
router bgp 65000
address-family l2vpn evpn
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 route-map rm_evpn_in in
exit-address-family
```

4. Configure the RR for EVPN.

a. Enable EVPN under the BGP configuration.

b. Configure both PE1 and PE3 as route-reflector clients.

c. Activate the EVPN address family for both peers.

```
router bgp 65000
bgp router-id 3.3.3.3
address-family l2vpn evpn
```

```

neighbor 1.1.1.1 activate
neighbor 1.1.1.1 route-reflector-client
neighbor 5.5.5.5 activate
neighbor 5.5.5.5 route-reflector-client
exit-address-family

```

Validation

- Verify BGP Neighbor Status.

```

PE1#show ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 3
1 BGP AS-PATH entries
3 BGP community entries

```

Neighbor Desc	V	AS	MsgRcv	MsgSen	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.3.3.3	4	65000	1312	1302	3	0	0	04:59:49	2

Total number of neighbors 1

Total number of Established sessions 1

- Verify that all BGP EVPN neighbor sessions are established.

```

show bgp l2vpn evpn summary
BGP router identifier 1.1.1.1, local AS number 65000
BGP table version is 15
1 BGP AS-PATH entries
3 BGP community entries

```

Neighbor AD MACIP MCAST	V	AS ESI	MsgRcv PREFIX-ROUTE	MsgSen Desc	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
3.3.3.3	4								
65000	1312	1302	15	0	0	04:59:49		5 1	3 1
0 0									

Total number of neighbors 1

Total number of Established sessions 1

- Verify that the route-maps are correctly applied.

```

PE1#show route-map
route-map rm_lul, permit, sequence 10
  Match clauses:
    ip address prefix-list: lu_pref1
  Set clauses:
    local-preference 300
route-map rm_lul, deny, sequence 11
  Match clauses:
    ip address prefix-list: lu_pref2
  Set clauses:
    originator-id 15.15.15.15
route-map rm_evpn_out, permit, sequence 10
  Match clauses:
    evpn-route-type: type-1
  Set clauses:
    local-preference 200
    community 65000:100 (additive)
route-map rm_evpn_out, permit, sequence 20
  Match clauses:
    evpn-route-type: type-2
    mac address list: mac_list1
  Set clauses:
    local-preference 150
    metric 50

```

```

community 65000:200 (additive)
route-map rm_evpn_out, permit, sequence 30
Match clauses:
  evpn-route-type: type-3
Set clauses:
  local-preference 100
  origin igp
community 65000:300 (additive)
route-map rm_evpn_out, permit, sequence 1000
Match clauses:
Set clauses:

```

- Verify EVPN-MPLS tunnel status.

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id Ext-Color FAT
=====
1.1.1.1          5.5.5.5          Installed    04:59:49     04:59:49     501          133501
---
Total number of entries are 1

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source          Destination      Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color FAT
=====
1.1.1.1          5.5.5.5          Installed    04:59:49     04:59:49     205          ---
---
Total number of entries are 1

```

- Verify EVPN MPLS Interfaces.

```

PE1#show evpn mpls
EVPN-MPLS Information
=====
Codes: NW - Network Port
       AC - Access Port
       (u) - Untagged

VPN-ID  EVI-Name  EVI-Type Type Interface ESI          VLAN  DF-Status
Src-Addr Dst-Addr
-----
205     ----    L2      NW    ----    ----          ----  ----
      1.1.1.1      5.5.5.5
205     ----    --      AC    xe25.205 --- Single Homed Port ---  ----  ----  -
---
Total number of entries are 2

Note: Refer sub-interface config for VLAN information.

```

- Verify BGP EVPN routes.

```

PE1#show bgp l2vpn evpn
BGP table version is 15, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

```


Description : Ext-Color - Extended community color

[EVPN route type]:[ESI]:[VNID]:[relevant route information]

- 1 - Ethernet Auto-discovery Route
- 2 - MAC/IP Route
- 3 - Inclusive Multicast Route
- 4 - Ethernet Segment Route
- 5 - Prefix Route

Hop	Network	Next Metric	LocPrf	Weight	Path	Peer	Encap
RD[1.1.1.1:105] VRF[vrf105]							
*>	[2]:[0]:[205]:[48,0011:2233:4455]:[32,98.98.101.1]:[17]	1.1.1.1		0	100	32768	i - -----
	MPLS						
*>	[2]:[0]:[205]:[48,0011:2233:4455]:[32,104.104.103.1]:[17]	1.1.1.1		0	100	32768	i - -----
	MPLS						
* i 1	[2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
* i 1	[2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
*>	[2]:[0]:[205]:[48,6821:5f1f:5220]:[32,103.103.102.1]:[17]	1.1.1.1		0	100	32768	i - -----
	MPLS						
* i 1	[2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
*>	[3]:[205]:[32,1.1.1.1]	1.1.1.1		0	100	32768	i - -----
	MPLS						
* i	[3]:[205]:[32,5.5.5.5]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
RD[1.1.1.1:501] VRF[Eline501]							
*>	[1]:[0]:[501]:[18]	1.1.1.1		0	100	32768	i - -----
	MPLS						
* i 1	[1]:[0]:[133501]:[18]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
RD[5.5.5.5:105]							
*> i 1	[2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
*> i 1	[2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
*> i 1	[2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
*> i	[3]:[205]:[32,5.5.5.5]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
RD[5.5.5.5:501]							
*> i 1	[1]:[0]:[133501]:[18]	5.5.5.5		0	100	0	i -
	3.3.3.3 MPLS						
Total number of prefixes 15							

- Verify MAC/IP table.

```

PE1#show bgp l2vpn evpn mac-ip
BGP table version is 15, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
                l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

```

```
RD[1.1.1.1:105] VRF[vrf105]:
```

ESI Address	VNID/LABEL	L3VNID	Eth-Tag Nexthop	Mac-Address GW-Type	IP-	Peer	Encap
*> 0			205	0011:2233:4455			
98.98.101.1	17	0	1.1.1.1	--		-----	MPLS
*> 0			205	0011:2233:4455			
104.104.103.1	17	0	1.1.1.1	--		-----	MPLS
* il 0			205	0011:2233:4567			
99.99.101.1	17	0	5.5.5.5	--		3.3.3.3	MPLS
* il 0			205	0011:2233:4567			
104.104.104.1	17	0	5.5.5.5	--		3.3.3.3	MPLS
*> 0			205	6821:5f1f:5220			
103.103.102.1	17	0	1.1.1.1	--		-----	MPLS
* il 0			205	e8c5:7a69:45ed			
103.103.103.1	17	0	5.5.5.5	--		3.3.3.3	MPLS

```
RD[5.5.5.5:105]
```

ESI Address	VNID/LABEL	L3VNID	Eth-Tag Nexthop	Mac-Address GW-Type	IP-	Peer	Encap
*>il 0			205	0011:2233:4567			
99.99.101.1	17	0	5.5.5.5	--		3.3.3.3	MPLS
*>il 0			205	0011:2233:4567			
104.104.104.1	17	0	5.5.5.5	--		3.3.3.3	MPLS
*>il 0			205	e8c5:7a69:45ed			
103.103.103.1	17	0	5.5.5.5	--		3.3.3.3	MPLS

```
PE1#show evpn mpls mac-table
```

```

=====
=====
=====
EVPN MPLS MAC Entries
=====
=====
=====
VNID      Interface VlanId   In-VlanId Mac-Addr      VTEP-
Ip/ESI    LeafFlag          Type        Status        MAC move AccessPortDesc
=====
```

205	irb305	----	----	0011.2233.4455				
1.1.1.1				Static Local	-----	0	-----	
205	----	----	----	0011.2233.4567				
5.5.5.5				Static Remote	-----	0	-----	
205	irb305	----	----	6821.5f1f.5220				
1.1.1.1				Static Local	-----	0	-----	
205	----	----	----	e8c5.7a69.45ed				
5.5.5.5				Static Remote	-----	0	-----	

```
Total number of entries are : 4
```

```
PE1#show evpn mpls arp-cache
MPLS-EVPN ARP-CACHE Information
=====
```

EVPN-ID	Ip-Addr	Mac-Addr	Type	Age-Out	Retries-Left
205	98.98.101.1	0011.2233.4455	Static Local	----	
205	99.99.101.1	0011.2233.4567	Static Remote	----	
205	103.103.102.1	6821.5f1f.5220	Static Local	----	

```

205      103.103.103.1      e8c5.7a69.45ed Static Remote ----
205      104.104.103.1      0011.2233.4455 Static Local ----
205      104.104.104.1      0011.2233.4567 Static Remote ----
Total number of entries are 6

PE1#show bgp l2vpn evpn detail
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4455]:[32,98.98.101.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4455]:[32,104.104.103.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,6821:5f1f:5220]:[32,103.103.102.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local

```

```
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [3]:[205]:[32,1.1.1.1]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 13:46:13 2025
Peer : Local
BGP route entry for prefix : [3]:[205]:[32,5.5.5.5]
Route-Distinguisher : [1.1.1.1:105] VRF : vrf105
Flags : Valid, IBGP
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [1]:[0]:[501]:[18]
Route-Distinguisher : [1.1.1.1:501] VRF : Eline501

Flags : Valid, Selected
Nexthop : 1.1.1.1 MED value: 0
Community :
Extended Community: RT:501:501 Encapsulation:MPLS ESI-Label:0 Control-Flags,MTU:2,9216
Weight : 32768, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 15:57:57 2025
Peer : Local
BGP route entry for prefix : [1]:[0]:[133501]:[18]
Route-Distinguisher : [1.1.1.1:501] VRF : Eline501
Flags : Valid, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:501:501 Encapsulation:MPLS ESI-Label:0 Control-Flags,MTU:2,9216
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
```

```

Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS MAC_mob_seq:Static
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
BGP route entry for prefix : [3]:[205]:[32,5.5.5.5]
Route-Distinguisher : [5.5.5.5:105]
Flags : Valid, Selected, IBGP
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:65000:1073742029 Encapsulation:MPLS
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025

Peer : 3.3.3.3
BGP route entry for prefix : [1]:[0]:[133501]:[18]
Route-Distinguisher : [5.5.5.5:501]
Flags : Valid, Selected, IBGP, Labelled, Labelled
Nexthop : 5.5.5.5 MED value : 0
Community :
Extended Community: RT:501:501 Encapsulation:MPLS ESI-Label:0 Control-Flags,MTU:2,9216
Weight : 0, Local Preference :100
AS Path : Local
Origin : IGP
Last Update : Wed Sep 24 17:36:32 2025
Peer : 3.3.3.3
Total number of prefixes 15
PE1#show bgp l2vpn evpn vrf vrf105
BGP table version is 1, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i
- internal,
l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color
[EVPN route type]:[ESI]:[VNID]:[relevent route informantion]
1 - Ethernet Auto-discovery Route
2 - MAC/IP Route
3 - Inclusive Multicast Route
4 - Ethernet Segment Route
5 - Prefix Route
Network Next
Hop Metric LocPrf Weight Path Peer Encap
*> [2]:[0]:[205]:[48,0011:2233:4455]:[32,98.98.101.1]:[17]
1.1.1.1 0 100 32768 i - -----
-- MPLS
*> [2]:[0]:[205]:[48,0011:2233:4455]:[32,104.104.103.1]:[17]
1.1.1.1 0 100 32768 i - -----
-- MPLS
* i 1 [2]:[0]:[205]:[48,0011:2233:4567]:[32,99.99.101.1]:[17]
5.5.5.5 0 100 0 i -
3.3.3.3 MPLS
* i 1 [2]:[0]:[205]:[48,0011:2233:4567]:[32,104.104.104.1]:[17]
5.5.5.5 0 100 0 i -
3.3.3.3 MPLS
*> [2]:[0]:[205]:[48,6821:5f1f:5220]:[32,103.103.102.1]:[17]
1.1.1.1 0 100 32768 i - -----
-- MPLS
* i 1 [2]:[0]:[205]:[48,e8c5:7a69:45ed]:[32,103.103.103.1]:[17]

```

```
5.5.5.5 0 100 0 i -  
3.3.3.3 MPLS  
*> [3]:[205]:[32,1.1.1.1]  
1.1.1.1 0 100 32768 i - -----  
-- MPLS  
* i [3]:[205]:[32,5.5.5.5]  
5.5.5.5 0 100 0 i -  
3.3.3.3 MPLS  
Total number of prefixes 8
```

Implementation Examples

- **EVPN Route Control Between Leaf and Spine Nodes** : Leaf nodes may advertise routes that other spines do not require. You can apply a route-map to filter those routes based on AS numbers or route types.
- **Community-Based Route Filtering**: When VRF routes are redistributed into EVPN, you can tag them with specific communities and use route-maps to advertise only selected routes.

Commands

The EVPN AF route-maps and route filtering feature introduces the following commands:

mac-list

Use this command to create a new MAC list that specifies MAC addresses with an associated mask and defines an action (permit or deny).

Use the **no** parameter of this command to unconfigure.

Command Syntax

Mac List:

```
mac-list WORD  
  
no mac-list WORD
```

Mac list entry:

```
permit/deny XXXX.XXXX.XXXX XXXX.XXXX.XXXX  
  
no permit/deny XXXX.XXXX.XXXX XXXX.XXXX.XXXX
```

Parameters

mac-list

Build a MAC list

WORD

Name of MAC list

Permit

Permits the MAC Address (XXXX.XXXX.XXXX) in HHHH.HHHH.HHHH format

Deny

Denies the MAC Address (XXXX.XXXX.XXXX) in HHHH.HHHH.HHHH format

Default

None

Command Mode

Configure mode and Mac-list mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to specifies MAC addresses with an associated mask and defines an action (permit or deny):

```
ocnos#configure terminal  
ocnos(config)#mac-list mlist1  
ocnos(config-mac-list)permit 1111.1111.3322 1100.0011.0000
```

match mac address list

Use this command to match a mac address list.

Use the **no** parameter of this command to turn off the matching.

Command Syntax

```
match mac address list WORD
no match mac address list WORD
```

Parameters

WORD

Defines the MAC list name

Default

None

Command Mode

Router Map mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to match a mac address list:

```
R1#configure terminal
R1(config)#route-map rmap1 permit 10
R1(config-route-map)# match mac address list mlist1
```


match evpn-route-type

Use this command to match an EVPN route by its route type. The supported route types are: **type-1**, **type-2**, **type-3**, **type-4**, **type-5**, **type-2-MAC-IP**, and **type-2-MAC-ONLY**.

Use the **no** parameter of this command to turn off the matching.

Command Syntax

```
match evpn-route-type (type-1|type-2|type-3|type-4|type-5|type-2-MAC-IP|type-2-MAC-ONLY|)
no match evpn-route-type (type-1|type-2|type-3|type-4|type-5|type-2-MAC-IP|type-2-MAC-ONLY|)
```

Parameters

type-1

Matches Ethernet Auto-Discovery (AD) routes

type-2

Matches both MAC-only and MAC/IP advertisement routes

type-2-MAC-IP

Matches MAC/IP advertisement routes

type-2-MAC-ONLY

Matches MAC-only advertisement routes

type-3

Matches Multicast Ethernet Tag (MET) routes

type-4

Matches Ethernet Segment Identifier (ESI) routes

type-5

Matches IP Prefix routes

Default

None

Command Mode

Router Map mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to match an EVPN route by its route type:

```
R1#configure terminal
R1(config)#route-map rmap1 permit 10
R1(config-route-map)# match evpn-route-type type-2
```

show mac-list

Use this command to view the detailed information about a specific mac-list.

Command Syntax

```
show mac-list detail WORD
```

Parameters

WORD

Specifies the name of the MAC list for which detailed information is to be displayed.

Default

None

Command Mode

Privileged execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following example illustrates how to view the detailed information about a specific mac-list:

```
ocnos#show mac-list detail mlist1
mac-list mlist1:
  count: 2, sequences: 3 - 5
  ripd:
  ripngd:
  ospfd:
  ospf6d:
  ldpd:
  bgpd:
  seq 3 permit 0000.aaaa.bbbb ffff.ffff.ffff (hit count: 10, refcount: 21)
  seq 5 deny 0000.ccaa.bbbb ffff.ffff.ffff (hit count: 1, refcount: 2)
```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms/Acronym	Description
Address Family Identifier / Subsequent AFI (AFI/SAFI)	Specifies network layer protocol and route type.
Border Gateway Protocol (BGP)	Routing protocol used to exchange routing information across autonomous systems.
Ethernet VPN (EVPN)	BGP-based control plane for Ethernet multipoint services.

BGP Labeled Unicast Next Hop Self in Route-Map

Overview

The BGP-LU next-hop-self in route map feature provides the ability set the local BGP peer as the next-hop for select BGP-LU routes. When such a route map is applied to a BGP-LU neighbor in the outbound direction, the matched routes will be updated as below:

- The next-hop address is replaced with the local BGP peer address. Based on the BGP configuration, this is either the local interface address or the local loopback address
- The label is replaced with the local label allocated for the prefix

Feature Characteristics

BGP-LU routes permitted by the route-map and configured with `set ip next-hop self` are advertised with the local BGP peer address as the next-hop and the locally assigned label. BGP-LU routes permitted by the route-map without `set next-hop self` are advertised with their original next-hop and the label remains unchanged. The BGP-LU routes denied by the route-map are not advertised.

Limitations

The existing route-map CLIs to set next-hop, such as, `set ip next-hop a.b.c.d` or `set ip next-hop peer-address` are not recommended to be used for this feature and they can impact adversely.

Prerequisites

While deploying BGP labelled unicast, ensure that the `allocate-label` command is enabled under `router bgp` mode. This command is required for allocating labels to IPv4 prefixes.

Configuration

Topology

This topology contains Route Reflectors (RR) and PE nodes with BGP-LU as the transport between them.

Domain 1 contains LDP as transport with OSPF as IGP on all nodes. Domain 2 contains RSVP as transport with IS-IS as IGP on all nodes.

Figure 4. BGP LU Next-hop in Route-map

To configure BGP LU Next-hop, follow the steps mentioned below:

1. Configure BGP LU as transport

```
#configure terminal
(config)#interface lo
(config-if)#ip address 11.11.11.55/32 secondary
(config-if)#exit
(config)#interface xe16
(config-if)#ip address 172.4.5.55/24
(config-if)#label-switching
(config-if)#exit
(config)#commit
```

2. Configure the routing process OSPF with process ID 1.

```
(config)#router ospf 1
```

3. Define the interface (172.4.5.0/24) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).

```
(config-router)#network 172.4.5.0/24 area 0
```

4. Define the interface (11.11.11.55/32) on which OSPF runs, and associate the area ID (0) with the interface (area ID 0 specifies the backbone area).

```
(config-router)#network 11.11.11.55/32 area 0
(config-router)#exit
(config)#commit
(config)#router bgp 100
```

5. Create a Prefix-list to match the advertised prefixes

```
(config-router)# ip prefix-list BGP-LU-FILTER permit 10.10.10.0/24
```

6. Create Route-map for LU using the prefix-list

```
(config-router)#route-map RM-BGPLU-OUT permit 10
(config-router)#match ip address prefix-list BGP-LU-FILTER
```

Optionally, apply strict filtering using deny

```
(config-router)#route-map RM-BGPLU-OUT deny 20
```

7. Apply Route-map on BGP-LU Neighbor

```
(config-router)# allocate-label all route-map RM-BGPLU-OUT
```

Configuration Snapshot

PE1

```
PE1#sh run
!
service password-encryption
!
logging console 3
logging monitor 5
logging logfile device_debug_log 2
logging level nsm 5
logging level ospf 5
logging level ldp 5
logging level hsl 5
logging level bgp 5
logging level cml 5
logging level cmm 4
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
forwarding profile kaps profile-two
hardware-profile filter qos enable
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE1
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
feature rsyslog
logging remote server 10.16.58.70 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
ip vrf management
!
ip vrf VRF1
  rd 100:1
  route-target both 100:1
!
ip vrf VRF2
```

```
rd 100:2
route-target both 100:2
!
router ldp
router-id 1.1.1.1
session-protection
pw-status-tlv
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
interface ge1
!
interface ge2
!
interface ge3
!
interface ge4
!
interface ge5
!
interface ge6
!
interface ge7
!
interface ge8
!
interface ge9
!
interface ge10
!
interface ge11
!
interface ge12
!
interface ge13
!
interface ge14
!
interface ge15
!
interface ge16
!
interface ge17
!
interface ge18
!
interface ge19
!
interface ge20
!
interface ge21
!
interface ge22
!
interface ge29
!
interface lo
ip address 127.0.0.1/8
ip address 1.1.1.1/32 secondary
ipv6 address ::1/128
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
```

```
interface xe23
!
interface xe24
!
interface xe25
!
interface xe26
  load-interval 30
  ip address 10.1.1.1/24
  mtu 9216
  label-switching
  ip ospf network point-to-point
  enable-ldp ipv4
!
interface xe27
  load-interval 30
  mtu 9216
!
interface xe27.102
  description L3VPN-VRF1
  encapsulation dot1q 102
  load-interval 30
  ip vrf forwarding VRF1
  ip address 101.1.1.1/24
  isis network point-to-point
  ip router isis 100
!
interface xe27.103
  description L3VPN-VRF2
  encapsulation dot1q 103
  load-interval 30
  ip vrf forwarding VRF2
  ip address 101.1.2.1/24
  isis network point-to-point
  ip router isis 200
!
interface xe28
!
  exit
!
router ospf 65535
  ospf router-id 1.1.1.1
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 10.1.1.0/24 area 0.0.0.0
!
router isis 100 VRF1
  is-type level-1-2
  metric-style wide
  dynamic-hostname
  bfd all-interfaces
  net 49.0001.0000.0000.0001.00
  redistribute bgp
!
router isis 200 VRF2
  is-type level-1-2
  metric-style wide
  dynamic-hostname
  bfd all-interfaces
  net 49.0002.0000.0000.0002.00
  redistribute bgp
!
router bgp 4200000001
  bgp router-id 1.1.1.1
  bgp auto-policy-soft-reset enable
  bgp log-neighbor-changes
  no bgp inbound-route-filter
  allocate-label all
```

```

neighbor 2.2.2.2 remote-as 4200000001
neighbor 2.2.2.2 tcp-mss 1440
neighbor 2.2.2.2 update-source 1.1.1.1
neighbor 2.2.2.2 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor 2.2.2.2 advertisement-interval 0
neighbor 2.2.2.2 fall-over bfd multihop
!
address-family ipv4 unicast
network 1.1.1.1/32
exit-address-family
!
address-family ipv4 labeled-unicast
neighbor 2.2.2.2 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor 2.2.2.2 activate
exit-address-family
!
address-family ipv4 vrf VRF1
redistribute connected
redistribute isis
exit-address-family
!
address-family ipv4 vrf VRF2
redistribute connected
redistribute isis
exit-address-family
!
exit
!
line console 0
  exec-timeout 0 0
line vty 0 16
  exec-timeout 0 0
!
!
end

```

PE2

```

PE2-7001#sh running-config
!
! Software version: UFI_S9600-56DX-OcNOS-SP-PLUS-7.0.0.168-Alpha
11/13/202
5 18:32:31
!
! Last configuration change at 11:31:56 UTC Tue Nov 18 2025 by root
!
feature netconf-ssh
feature netconf-tls
!
feature netconf notification-cache enable
max-cache-notifications 0
!
cml bulk-config limit cpu enable
!
background-debug
  log all
  level 7
  suppress-non-bdr-logs
  exit
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps pwdelete
snmp-server enable traps pw

```



```
snmp-server enable traps mpls
snmp-server enable traps snmp authentication
snmp-server enable traps ospf
snmp-server enable traps bgp
snmp-server enable traps isis
!
bgp extended-asn-cap
!
hardware-profile filter evpn-mpls-mh enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
hardware-profile port-config mode3
!
bfd interval 3 minrx 3 multiplier 3
!
key chain isis
  key-id 3
    key-string encrypted 0xc8a471564ac751dc
!
key chain BGP
  key-id 4
    key-string encrypted
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157b
                                e7f6e35f32bf637
!
qos enable
qos statistics
!
mpls ilm-ecmp ldp
mpls ftn-ecmp ldp
mpls label mode vpnv4 all-vrfs per-vrf
mpls label mode vpnv6 all-vrfs per-vrf
!
mpls vpls vpls4294961250 1250
  control-word
  signaling bgp
  ve-id 61250
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961251 1251
  control-word
  signaling bgp
  ve-id 61251
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961252 1252
  control-word
  signaling bgp
  ve-id 61252
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961253 1253
  control-word
  signaling bgp
  ve-id 61253
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961254 1254
  control-word
  signaling bgp
  ve-id 61254
  exit-signaling
  exit-vpls
!
```

```
mpls vpls vpls4294961255 1255
  control-word
  signaling bgp
  ve-id 61255
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961256 1256
  control-word
  signaling bgp
  ve-id 61256
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961257 1257
  control-word
  signaling bgp
  ve-id 61257
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961258 1258
  control-word
  signaling bgp
  ve-id 61258
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961259 1259
  control-word
  signaling bgp
  ve-id 61259
  exit-signaling
  exit-vpls
!
mpls vpls vpls4294961260 1260
  control-word
  signaling bgp
  ve-id 61260
  exit-signaling
  exit-vpls
!
mpls vpls VPLS_AD1301 61301
  vpls-mtu 8000
  control-word
  signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
  l2vpn-id 192.168.36.2:1301
  rd 4294967294:61301
  route-target both 4294967294:61301
  exit-bgp-auto-discovery
  exit-signaling
  exit-vpls
!
mpls vpls vpls1261 4294961261
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1262 4294961262
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
```

```
    vpls-peer 192.168.36.3
    vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1263 4294961263
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1264 4294961264
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1265 4294961265
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1266 4294961266
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1267 4294961267
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1268 4294961268
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls vpls1269 4294961269
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
```

```
mpls vpls vpls1270 4294961270
  control-word
  signaling ldp
  vpls-peer 192.168.36.1
  vpls-peer 192.168.36.3
  vpls-peer 192.168.36.4
  exit-signaling
  exit-vpls
!
mpls vpls VPLS_AD1302 4294961302
  vpls-mtu 8000
  control-word
  signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961302
    rd 4294967294:61302
    route-target both 4294967294:61302
  exit-bgp-auto-discovery
  exit-signaling
  exit-vpls
!
mpls vpls VPLS_AD1303 4294961303
  vpls-mtu 8000
  control-word
  signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961303
    rd 4294967294:61303
    route-target both 4294967294:61303
  exit-bgp-auto-discovery
  exit-signaling
  exit-vpls
!
mpls vpls VPLS_AD1304 4294961304
  vpls-mtu 8000
  control-word
  signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961304
    rd 4294967294:61304
    route-target both 4294967294:61304
  exit-bgp-auto-discovery
  exit-signaling
  exit-vpls
!
mpls vpls VPLS_AD1305 4294961305
  vpls-mtu 8000
  control-word
  signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961305
    rd 4294967294:61305
    route-target both 4294967294:61305
  exit-bgp-auto-discovery
  exit-signaling
  exit-vpls
!
mpls vpls VPLS_AD1306 4294961306
  vpls-mtu 8000
  control-word
  signaling ldp
  vpls-type ethernet
  bgp-auto-discovery
    l2vpn-id 65000:4294961306
```

```
        rd 4294967294:61306
        route-target both 4294967294:61306
    exit-bgp-auto-discovery
    exit-signaling
    exit-vpls
    !
mpls vpls VPLS_AD1307 4294961307
    vpls-mtu 8000
    control-word
    signaling ldp
    vpls-type ethernet
    bgp-auto-discovery
        l2vpn-id 65000:4294961307
        rd 4294967294:61307
        route-target both 4294967294:61307
    exit-bgp-auto-discovery
    exit-signaling
    exit-vpls
    !
mpls vpls VPLS_AD1308 4294961308
    vpls-mtu 8000
    control-word
    signaling ldp
    vpls-type ethernet
    bgp-auto-discovery
        l2vpn-id 65000:4294961308
        rd 4294967294:61308
        route-target both 4294967294:61308
    exit-bgp-auto-discovery
    exit-signaling
    exit-vpls
    !
mpls vpls VPLS_AD1309 4294961309
    vpls-mtu 8000
    control-word
    signaling ldp
    vpls-type ethernet
    bgp-auto-discovery
        l2vpn-id 65000:4294961309
        rd 4294967294:61309
        route-target both 4294967294:61309
    exit-bgp-auto-discovery
    exit-signaling
    exit-vpls
    !
mpls vpls VPLS_AD1310 4294961310
    vpls-mtu 8000
    control-word
    signaling ldp
    vpls-type ethernet
    bgp-auto-discovery
        l2vpn-id 65000:4294961310
        rd 4294967294:61310
        route-target both 4294967294:61310
    exit-bgp-auto-discovery
    exit-signaling
    exit-vpls
    !
mpls l2-circuit PE2-To-PE1-1271 1271 192.168.36.1
    control-word
    !
mpls l2-circuit PE2-To-PE1-1272 4294961272 192.168.36.1
    control-word
    !
mpls l2-circuit PE2-To-PE1-1273 4294961273 192.168.36.1
    control-word
    !
mpls l2-circuit PE2-To-PE1-1274 4294961274 192.168.36.1
```

```
control-word
!
mpls 12-circuit PE2-To-PE1-1275 4294961275 192.168.36.1
control-word
!
mpls 12-circuit PE2-To-PE1-1276 4294961276 192.168.36.1
control-word
!
mpls 12-circuit PE2-To-PE1-1277 4294961277 192.168.36.1
control-word
!
mpls 12-circuit PE2-To-PE1-1278 4294961278 192.168.36.1
control-word
!
mpls 12-circuit PE2-To-PE1-1279 4294961279 192.168.36.1
control-word
!
mpls 12-circuit PE2-To-PE1-1280 4294961280 192.168.36.1
control-word
!
mpls 12-circuit PE2-To-PE3-1281 4294961281 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1282 4294961282 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1283 4294961283 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1284 4294961284 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1285 4294961285 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1286 4294961286 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1287 4294961287 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1288 4294961288 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1289 4294961289 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE3-1290 4294961290 192.168.36.3
control-word
!
mpls 12-circuit PE2-To-PE4-1291 4294961291 192.168.36.4
control-word
!
mpls 12-circuit PE2-To-PE4-1292 4294961292 192.168.36.4
control-word
!
mpls 12-circuit PE2-To-PE4-1293 4294961293 192.168.36.4
control-word
!
mpls 12-circuit PE2-To-PE4-1294 4294961294 192.168.36.4
control-word
!
mpls 12-circuit PE2-To-PE4-1295 4294961295 192.168.36.4
control-word
!
mpls 12-circuit PE2-To-PE4-1296 4294961296 192.168.36.4
control-word
!
```

```
mpls 12-circuit PE2-To-PE4-1297 4294961297 192.168.36.4
  control-word
!
mpls 12-circuit PE2-To-PE4-1298 4294961298 192.168.36.4
  control-word
!
mpls 12-circuit PE2-To-PE4-1299 4294961299 192.168.36.4
  control-word
!
mpls 12-circuit PE2-To-PE4-1300 4294961300 192.168.36.4
  control-word
!
mpls 12-circuit vc2000 4294967290 192.168.36.11
!
mpls 12-circuit vc2001 4294967290 192.168.36.12
!
mpls 12-circuit vc2004 4294967291 192.168.36.11
!
mpls 12-circuit vc2003 4294967291 192.168.36.12
!
hostname PE2-7001
port ce46 breakout 4X10g
port ce47 breakout 4X10g
ip name-server vrf management 10.12.3.23
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
ospf restart grace-period 2
ospf restart helper max-grace-period 2
aaa local authentication attempts max-fail 25
aaa local authentication unlock-timeout 1
aaa authentication login error-enable
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community public vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature rsyslog
logging remote server 10.16.100.20 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
router-id 192.168.36.2
!
evpn mpls enable
!
evpn mpls irb
!
ip vrf irbvrf3001
  rd 4294967294:63001
  route-target both 4294967294:63001
  l3vni 23001
  maximum-fib-routes ipv4 10000 warning-only
  maximum-fib-routes ipv6 10000 warning-only
!
ip vrf irbvrf3002
  rd 192.168.36.2:3002
  route-target import 4294967294:63002
  route-target export 4294967294:63022
  l3vni 23002
```

```
!  
ip vrf irbvrf3003  
  rd 192.168.36.2:3003  
  route-target import 4294967294:63003  
  route-target export 4294967294:63033  
  export map RM-EXPORT-EVPN-IRBVRF3003-IPv6  
  l3vni 23003  
!  
ip vrf management  
!  
ip vrf vrf101  
  rd 4294967294:101  
  route-target both 4294967294:101  
  maximum-fib-routes ipv4 10000 stop-install  
  maximum-fib-routes ipv6 10000 stop-install  
!  
ip vrf vrf102  
  rd 4294967294:100  
  route-target both 4294967294:102  
  export map RM-EXPORT-ATTR  
!  
ip vrf vrf103  
  rd 4294967294:103  
  route-target both 4294967294:103  
  export map RM-EXPORT-EVPN-IRBVRF3003-IPv6  
!  
ip vrf vrf104  
  rd 4294967294:104  
  route-target both 4294967294:104  
!  
ip vrf vrf105  
  rd 4294967294:105  
  route-target both 4294967294:105  
!  
ip vrf vrf106  
  rd 4294967294:106  
  route-target both 4294967294:106  
!  
ip vrf vrf107  
  rd 4294967294:107  
  route-target both 4294967294:107  
!  
ip vrf vrf108  
  rd 4294967294:108  
  route-target both 4294967294:108  
!  
ip vrf vrf109  
  rd 4294967294:109  
  route-target both 4294967294:109  
!  
ip vrf vrf110  
  rd 4294967294:11  
  route-target both 4294967294:110  
  maximum-fib-routes ipv4 10000 stop-install  
  maximum-fib-routes ipv6 10000 stop-install  
!  
ip vrf vrf111  
  rd 4294967294:111  
  route-target both 4294967294:111  
  export map RM-EXPORT-VRF111  
!  
mac vrf ELAN1901  
  rd 192.168.36.2:1901  
  route-target both 1901:1901  
  export map RM-EXPORT-ELAN1901  
!  
mac vrf ELAN1902  
  rd 4294967294:61902
```



```
    route-target both 4294967294:61902
!
mac vrf ELAN1903
  rd 4294967294:1903
  route-target both 4294967294:1903
!
mac vrf ELAN1904
  rd 4294967294:61904
  route-target both 4294967294:41904
!
mac vrf ELAN1905
  rd 4294967294:1905
  route-target both 4294967294:21901
!
mac vrf ELAN1906
  rd 4294967294:1906
  route-target both 4294967294:41906
!
mac vrf ELAN1907
  rd 4294967294:1907
  route-target both 4294967294:41907
!
mac vrf ELAN1908
  rd 4294967294:1908
  route-target both 4294967294:41908
!
mac vrf ELAN1909
  rd 4294967294:1909
  route-target both 4294967294:41909
!
mac vrf ELAN1910
  rd 4294967294:1910
  route-target both 4294967294:41910
!
mac vrf eline901
  rd 192.168.36.2:901
  route-target both 901:901
!
mac vrf eline902
  rd 4294967294:902
  route-target both 4294967294:902
!
mac vrf eline903
  rd 4294967294:903
  route-target both evpn-auto-rt
!
mac vrf eline904
  rd 192.168.36.2:904
  route-target both evpn-auto-rt
  route-target both 904:904
!
mac vrf eline905
  rd 192.168.36.2:905
  route-target both 905:905
!
mac vrf eline906
  rd 192.168.36.2:906
  route-target both 906:906
!
mac vrf eline907
  rd 192.168.36.2:907
  route-target both 907:907
!
mac vrf eline908
  rd 192.168.36.2:908
  route-target both 908:908
!
mac vrf eline909
```

```
rd 192.168.36.2:909
route-target both 909:909
!
mac vrf eline910
rd 192.168.36.2:910
route-target both 910:910
!
mac vrf irbElan3001
rd 192.168.36.2:1001
route-target both 4294967294:3001
!
mac vrf irbElan3002
rd 192.168.36.2:1002
route-target export 4294967294:301
route-target import 4294967294:302
!
mac vrf irbElan3003
rd 192.168.36.2:1003
route-target both 4294967294:3003
!
evpn mpls vtep-ip-global 192.168.36.2
!
evpn mpls mac-ageing-time 20
!
evpn mpls id 1901
host-reachability-protocol evpn-bgp ELAN1901
!
evpn mpls id 61902
host-reachability-protocol evpn-bgp ELAN1902
!
evpn mpls id 61904
host-reachability-protocol evpn-bgp ELAN1904
!
evpn mpls id 163001
host-reachability-protocol evpn-bgp irbElan3001
evpn irb irb3001
!
evpn mpls id 163002
host-reachability-protocol evpn-bgp irbElan3002
evpn irb irb3002
!
evpn mpls id 163003
host-reachability-protocol evpn-bgp irbElan3003
evpn irb irb3003
!
evpn mpls id 1671905
host-reachability-protocol evpn-bgp ELAN1905
!
evpn mpls id 1671906
host-reachability-protocol evpn-bgp ELAN1906
!
evpn mpls id 1671907
host-reachability-protocol evpn-bgp ELAN1907
!
evpn mpls id 1671908
host-reachability-protocol evpn-bgp ELAN1908
!
evpn mpls id 1671909
host-reachability-protocol evpn-bgp ELAN1909
!
evpn mpls id 1671910
host-reachability-protocol evpn-bgp ELAN1910
!
evpn mpls id 16772901 xconnect target-mpls-id 16771901
host-reachability-protocol evpn-bgp eline901
!
evpn mpls id 16772902 xconnect target-mpls-id 16771902
host-reachability-protocol evpn-bgp eline902
```

```
!  
evpn mpls id 16772903 xconnect target-mpls-id 16771903  
  host-reachability-protocol evpn-bgp eline903  
!  
evpn mpls id 16772904 xconnect target-mpls-id 16771904  
  host-reachability-protocol evpn-bgp eline904  
!  
evpn mpls id 16772905 xconnect target-mpls-id 16771905  
  host-reachability-protocol evpn-bgp eline905  
!  
evpn mpls id 16772906 xconnect target-mpls-id 16771906  
  host-reachability-protocol evpn-bgp eline906  
!  
evpn mpls id 16772907 xconnect target-mpls-id 16771907  
  host-reachability-protocol evpn-bgp eline907  
!  
evpn mpls id 16772908 xconnect target-mpls-id 16771908  
  host-reachability-protocol evpn-bgp eline908  
!  
evpn mpls id 16772909 xconnect target-mpls-id 16771909  
  host-reachability-protocol evpn-bgp eline909  
!  
evpn mpls id 16772910 xconnect target-mpls-id 16771910  
  host-reachability-protocol evpn-bgp eline910  
!  
evpn mpls id 16777215  
  host-reachability-protocol evpn-bgp ELAN1903  
!  
segment-routing  
!  
ip multicast-routing  
!  
ipv6 multicast-routing  
!  
ip prefix-list DEFAULT  
  seq 10 permit 0.0.0.0/0  
!  
ip prefix-list LDP  
  seq 5 deny any  
!  
ip prefix-list LOOPBACK  
  seq 10 permit 192.168.36.2/32  
!  
ip prefix-list PL-CUST-SUBNETS-IRBVRF3003-IPv4  
  seq 5 permit 192.2.2.0/24  
!  
ip prefix-list PL-DENY-DEFAULT-IRBVRF3003-IPv4  
  seq 5 permit 0.0.0.0/0  
!  
ip prefix-list PL-DNS-SERVERS-IRBVRF3003-IPv4  
  seq 5 permit 172.16.30.53/32  
!  
ip prefix-list PL-EXPORT-INTERVRFv4  
  seq 5 deny 5.5.5.0/24  
  seq 10 permit 5.5.6.0/24  
  seq 11 permit 5.5.7.0/24  
!  
ip prefix-list PL-NEXTHOP-IRBVRF3003-IPv4  
  seq 5 permit 80.12.1.254/32  
  seq 10 permit 201.103.1.2/32  
!  
ip prefix-list PL-NOMETRIC-IRBVRF3003-IPv4  
  seq 5 permit 192.168.30.0/24  
!  
mac-list PL-ELAN1901-MAC-HOSTS  
  seq 10 permit 0010.9400.0002 0010.9400.0002  
!  
ipv6 prefix-list PL-CUST-SUBNETS-IRBVRF3003-IPv6
```

```
seq 5 permit 2001:db8:3003:20::/64
!
ipv6 prefix-list PL-DENY-DEFAULT-IRBVR3003-IPv6
seq 5 permit ::/0
!
ipv6 prefix-list PL-DNS-SERVERS-IRBVR3003-IPv6
seq 5 permit 2001:db8:3003:53::53/128
!
ipv6 prefix-list PL-NEXTHOP-IRBVR3003-IPv6
seq 5 permit 80:12:1::254/128
seq 10 permit 2001:bd8:103::1/128
!
ipv6 prefix-list PL-NOMETRIC-IRBVR3003-IPv6
seq 5 permit 2001:db8:3003:10::/64
!
ipv6 prefix-list PL-VRFLAKINGv6
seq 5 deny 2222:1:1:1::/64
seq 6 permit 2222:1:1:2::/64
seq 7 permit 2222:1:1:3::/64
!
router ldp
router-id 192.168.36.2
fast-reroute
pw-status-tlv
ignore-mac-withdraw-bad-pdu-length
targeted-peer ipv4 192.168.36.1
exit-targeted-peer-mode
targeted-peer ipv4 192.168.36.3
exit-targeted-peer-mode
targeted-peer ipv4 192.168.36.4
exit-targeted-peer-mode
targeted-peer ipv4 192.168.36.11
exit-targeted-peer-mode
targeted-peer ipv4 192.168.36.12
exit-targeted-peer-mode
transport-address ipv4 192.168.36.2
neighbor all tcp-mss 1440
neighbor all auth md5 password plain-text P@ssw0rd
!
router rsvp
lsp-reoptimization-timer 2
hello-interval 3
hello-timeout 11
from 192.168.36.2
detour-allow-primary-upstream-path
detour-identification path
entropy-label-capability
auto-bypass
attributes best-effort
reoptimize
exit
inactivity-timer 30
enable
exit
auto-bandwidth-on-boot 1 5 1
!
route-map REDISTRIBUTE-CONNECTED-TO-BGP permit 10
match ip address prefix-list LOOPBACK
!
route-map RM-EXPORT-ELAN1901 permit 710
match mac address list PL-ELAN1901-MAC-HOSTS
set metric 300
set local-preference 300
set aigp-metric 100
set atomic-aggregate
set community 100:111
set large-community 1:1:1
set extcommunity rt 1:1 additive
```

```
!  
route-map RM-EXPORT-ATTR permit 10  
  match ip address prefix-list PL-EXPORT-INTERVRFv4  
  set tag 1000  
  set metric 4294967295  
  set local-preference 4294967295  
  set origin igp  
  set aigp-metric 2000  
  set as-path tag  
  set community 100:1 additive  
  set large-community 22:22:22  
  set extcommunity rt 120:10 230:10 additive  
!  
route-map RM-EXPORT-ATTR permit 20  
  match ipv6 address prefix-list PL-VRFLEAKINGv6  
  set tag 1001  
  set metric 4294967295  
  set local-preference 4294967295  
  set origin igp  
  set aigp-metric 3000  
  set as-path tag  
  set community 100:1 additive  
  set large-community 23:34:45  
  set extcommunity rt 120:120 230:130 additive  
!  
route-map RM-EXPORT-ATTR permit 30  
!  
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1001  
  match ip address prefix-list PL-NOMETRIC-IRBVRF3003-IPv4  
  continue 1002  
  set metric +1001  
  set local-preference 1001  
  set aigp-metric 1001  
  set extcommunity rt 51185:1001 additive  
!  
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1002  
  match ip address prefix-list PL-CUST-SUBNETS-IRBVRF3003-IPv4  
  continue 1003  
  set metric +1002  
  set local-preference 1002  
  set community 51185:1002 additive  
  set large-community 4294967294:51185:1002 additive  
!  
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1003  
  match ip address prefix-list PL-DENY-DEFAULT-IRBVRF3003-IPv4  
  continue 1004  
  set metric +1003  
  set local-preference 1003  
  set aigp-metric 1003  
  set extcommunity rt 51185:1003 additive  
!  
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1004  
  match ip address prefix-list PL-DNS-SERVERS-IRBVRF3003-IPv4  
  continue 1005  
  set tag 1001  
  set extcommunity color 4200000001  
  set atomic-aggregate  
  set metric 2345  
  set local-preference 4567  
  set origin igp  
  set aigp-metric 3000  
  set as-path tag  
  set community 4200000002  
  set large-community 65000:400:40  
  set aggregator as 65080 192.0.2.80  
  set extcommunity rt 4200000001:20 additive  
  set extcommunity cost 99 900  
!
```

```
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1005
match ip next-hop prefix-list PL-NEXTHOP-IRBVRF3003-IPv4
continue 1006
set metric +1005
set local-preference 1005
set aigp-metric 1005
set extcommunity rt 51185:1005 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1006
match community CUSTOMER_ROUTES_3015_IRBVRF3003_IPv4
continue 1007
set metric +1006
set local-preference 1006
set community 51185:1006 additive
set large-community 4294967294:51185:1006 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1007
match extcommunity CUSTOMER_EXTENDED_COMM_3016_IRBVRF3003_IPv4
continue 1008
set metric +1007
set local-preference 1007
set aigp-metric 1007
set extcommunity rt 51185:1007 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1008
match extcommunity CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVRF3003_IPv4
continue 1009
set metric +1008
set local-preference 1008
set community 51185:1008 additive
set large-community 4294967294:51185:1008 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1009
match large-community CUSTOMER_LARGE_COMM_IRBVRF3003_V4
continue 1010
set metric +1009
set local-preference 1009
set aigp-metric 1009
set extcommunity rt 51185:1009 additive
!
route-map RM-EXPORT-EVPN-IRBVRF3003-IPv6 permit 1010
match as-path ASPATH-IRBVRF3003-IN-V4
continue 1011
set metric +1010
set local-preference 1010
set community 51185:1010 additive
set large-community 4294967294:51185:1010 additive
!
interface cd48
!
interface cd49
!
interface cd50
!
interface cd51
!
interface cd52
!
interface cd53
!
interface cd54
!
interface cd55
!
interface ce1
!
interface ce2
!
```

```
interface ce3
!  
interface ce4
!  
interface ce5
!  
interface ce6
!  
interface ce7
!  
interface ce8
!  
interface ce9
!  
interface ce10
!  
interface ce11
!  
interface ce12
!  
interface ce13
!  
interface ce14
!  
interface ce15
!  
interface ce16
!  
interface ce17
!  
interface ce18
!  
interface ce19
!  
interface ce20
!  
interface ce21
!  
interface ce22
!  
interface ce23
!  
interface ce24
!  
interface ce25
!  
interface ce26
!  
interface ce27
!  
interface ce28
!  
interface ce29
!  
interface ce30
!  
interface ce31
!  
interface ce32
!  
interface ce33
!  
interface ce34
!  
interface ce35
!  
interface ce36
!
```

```
interface ce37
!
interface ce38
!
interface ce39
!
interface ce40
!
interface ce41
!
interface ce42
!
interface ce43
!
interface ce44
!
interface ce45
!
interface ce46/1
!
interface ce46/2
description connected_to_PE3
load-interval 30
ip address 203.0.113.18/31
ipv6 address 203:3:8::105/64
mtu 9194
label-switching
link-debounce-time 2000 0
mpls ldp-igp sync isis level-2 holddown-timer 900
isis network point-to-point
ip router isis 1
ipv6 router isis 1
isis authentication mode md5 level-1
isis authentication mode md5 level-2
isis authentication key-chain isis level-1
isis authentication key-chain isis level-2
mpls ldp-igp sync-delay 30
enable-rsvp
ip pim sparse-mode
lldp-agent
set lldp enable txrx
set lldp chassis-id-tlv ip-address
set lldp port-id-tlv if-name
lldp tlv basic-mgmt system-name select
lldp tlv basic-mgmt system-description select
exit
bfd interval 3 minrx 3 multiplier 3
!
interface ce46/3
description ### Link to RR2 ##
load-interval 30
mtu 9198
link-debounce-time 2000 0
lldp-agent
set lldp enable txrx
set lldp chassis-id-tlv ip-address
set lldp port-id-tlv if-name
lldp tlv basic-mgmt system-name select
lldp tlv basic-mgmt system-description select
exit
!
interface ce46/3.101
description ### Link to rr-2 ##
encapsulation dot1q 101
load-interval 30
ip address 203.0.113.11/31
ipv6 address 203:3:6::105/64
mtu 9194
```



```
label-switching
mpls ldp-igp sync isis level-2 holddown-timer 900
isis network point-to-point
ip router isis 1
ipv6 router isis 1
isis authentication mode md5 level-1
isis authentication mode md5 level-2
isis authentication key-chain isis level-1
isis authentication key-chain isis level-2
enable-ldp ipv4
mpls ldp-igp sync-delay 30
enable-rsvp
ip pim sparse-mode
bfd interval 3 minrx 3 multiplier 3
!
interface ce46/4
description connected_to_rr1
load-interval 30
ip address 203.0.113.16/31
ipv6 address 203:3:7::105/64
mtu 9194
label-switching
link-debounce-time 2000 0
mpls ldp-igp sync isis level-2 holddown-timer 900
isis network point-to-point
ip router isis 1
ipv6 router isis 1
isis authentication mode md5 level-1
isis authentication mode md5 level-2
isis authentication key-chain isis level-1
isis authentication key-chain isis level-2
mpls ldp-igp sync-delay 30
enable-rsvp
ip pim sparse-mode
lldp-agent
set lldp enable txrx
set lldp chassis-id-tlv ip-address
set lldp port-id-tlv if-name
lldp tlv basic-mgmt system-name select
lldp tlv basic-mgmt system-description select
exit
bfd interval 3 minrx 3 multiplier 3
!
interface ce47/1
!
interface ce47/2
!
interface ce47/3
!
interface ce47/4
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
interface irb3001
ip vrf forwarding irbvrf3001
ip address 80.10.1.1/24
ipv6 address 80:10:1::1/64
mtu 9216
!
interface irb3002
ip vrf forwarding irbvrf3002
ip address 80.11.1.1/24
ipv6 address 80:11:1::1/64
mtu 9216
!
interface irb3003
```

```
ip vrf forwarding irbvrf3003
ip address 80.12.1.1/24
ipv6 address 80:12:1::1/64
mtu 9216
!
interface lo
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface loopback1
ip address 192.168.36.2/32
ipv6 address cafe:168:36::2/128
prefix-sid index 2 explicit-null
ip router isis 1
ipv6 router isis 1
ip pim sparse-mode
!
interface xe0
!
interface xe1
description ## Connected to Spirent-2/16 ##
load-interval 30
mtu 9216
!
interface xe1.10 switchport
encapsulation dot1q 10
load-interval 30
mtu 8000
access-if-vpws
mpls-l2-circuit vc2001 primary
mpls-l2-circuit vc2000 secondary
!
interface xe1.11 switchport
encapsulation dot1q 11
load-interval 30
mtu 8000
access-if-vpws
mpls-l2-circuit vc2003 primary
mpls-l2-circuit vc2004 secondary
!
interface xe1.101
encapsulation dot1q 101
load-interval 30
ip vrf forwarding vrf101
ip address 201.101.1.1/24
ipv6 address 2001:bd8:101::1/64
mtu 9216
!
interface xe1.102
encapsulation dot1q 102
load-interval 30
ip vrf forwarding vrf102
ip address 201.102.1.1/24
ipv6 address 2001:bd8:102::1/64
mtu 9216
!
interface xe1.103
encapsulation dot1q 103
load-interval 30
ip vrf forwarding vrf103
ip address 201.103.1.1/24
ipv6 address 2001:bd8:103::1/64
mtu 9216
```

```
!  
interface xel.104  
  encapsulation dot1q 104  
  load-interval 30  
  ip vrf forwarding vrf104  
  ip address 201.104.1.1/24  
  ipv6 address 2001:bd8:104::1/64  
  mtu 9216  
!  
interface xel.105  
  encapsulation dot1q 105  
  load-interval 30  
  ip vrf forwarding vrf105  
  ip address 201.105.1.1/24  
  ipv6 address 2001:bd8:105::1/64  
  mtu 9216  
!  
interface xel.106  
  encapsulation dot1q 106  
  load-interval 30  
  ip vrf forwarding vrf106  
  ip address 201.106.1.1/24  
  ipv6 address 2001:bd8:106::1/64  
  mtu 9216  
!  
interface xel.107  
  encapsulation dot1q 107  
  load-interval 30  
  ip vrf forwarding vrf107  
  ip address 201.107.1.1/24  
  ipv6 address 2001:bd8:107::1/64  
  mtu 9216  
!  
interface xel.108  
  encapsulation dot1q 108  
  load-interval 30  
  ip vrf forwarding vrf108  
  ip address 201.108.1.1/24  
  ipv6 address 2001:bd8:108::1/64  
  mtu 9216  
!  
interface xel.109  
  encapsulation dot1q 109  
  load-interval 30  
  ip vrf forwarding vrf109  
  ip address 201.109.1.1/24  
  ipv6 address 2001:bd8:109::1/64  
  mtu 9216  
!  
interface xel.110  
  encapsulation dot1q 110  
  load-interval 30  
  ip vrf forwarding vrf110  
  ip address 201.110.1.1/24  
  ipv6 address 2001:bd8:110::1/64  
  mtu 9216  
!  
interface xel.111  
  encapsulation dot1q 111  
  load-interval 30  
  ip vrf forwarding vrf111  
  ip address 201.111.1.1/24  
  ipv6 address 2001:bd8:111::1/64  
  mtu 9216  
!  
interface xel.112  
  encapsulation dot1q 112  
  load-interval 30
```

```
ip vrf forwarding vrf111
ip address 201.112.1.1/24
ipv6 address 2001:bd8:112::1/64
mtu 9216
!
interface xel.113
encapsulation dot1q 113
load-interval 30
ip vrf forwarding vrf111
ip address 201.113.1.1/24
ipv6 address 2001:bd8:113::1/64
mtu 9216
!
interface xel.114
encapsulation dot1q 114
load-interval 30
ip vrf forwarding vrf111
ip address 201.114.1.1/24
ipv6 address 2001:bd8:114::1/64
mtu 9216
!
interface xel.115
encapsulation dot1q 115
load-interval 30
ip vrf forwarding vrf111
ip address 201.115.1.1/24
ipv6 address 2001:bd8:115::1/64
mtu 9216
!
interface xel.116
encapsulation dot1q 116
load-interval 30
ip vrf forwarding vrf111
ip address 201.116.1.1/24
ipv6 address 2001:bd8:116::1/64
mtu 9216
!
interface xel.117
encapsulation dot1q 117
load-interval 30
ip vrf forwarding vrf111
ip address 201.117.1.1/24
ipv6 address 2001:bd8:117::1/64
mtu 9216
!
interface xel.118
encapsulation dot1q 118
load-interval 30
ip vrf forwarding vrf111
ip address 201.118.1.1/24
ipv6 address 2001:bd8:118::1/64
mtu 9216
!
interface xel.119
encapsulation dot1q 119
load-interval 30
ip vrf forwarding vrf111
ip address 201.119.1.1/24
ipv6 address 2001:bd8:119::1/64
mtu 9216
!
interface xel.120
encapsulation dot1q 120
load-interval 30
ip vrf forwarding vrf111
ip address 201.120.1.1/24
ipv6 address 2001:bd8:11a::1/64
mtu 9216
```

```
!  
interface xel.300  
!  
interface xel.890  
  description for ### IPv4 eBGP  
  encapsulation dot1q 890  
  load-interval 30  
  ip address 190.160.2.1/24  
  mtu 9216  
!  
interface xel.891  
  description for ### 6PE  
  encapsulation dot1q 891  
  load-interval 30  
  ipv6 address 3601::1/64  
  mtu 9216  
!  
interface xel.892  
  description for ### 6PE  
  encapsulation dot1q 892  
  load-interval 30  
  ipv6 address 3602::1/64  
  mtu 9216  
!  
interface xel.893  
  description for ### 6PE  
  encapsulation dot1q 893  
  load-interval 30  
  ipv6 address 3603::1/64  
  mtu 9216  
!  
interface xel.894  
  description for ### 6PE  
  encapsulation dot1q 894  
  load-interval 30  
  ipv6 address 3604::1/64  
  mtu 9216  
!  
interface xel.895  
  description for ### 6PE  
  encapsulation dot1q 895  
  load-interval 30  
  ipv6 address 3605::1/64  
  mtu 9216  
!  
interface xel.896  
  description for ### 6PE  
  encapsulation dot1q 896  
  load-interval 30  
  ipv6 address 3606::1/64  
  mtu 9216  
!  
interface xel.897  
  description for ### 6PE  
  encapsulation dot1q 897  
  load-interval 30  
  ipv6 address 3607::1/64  
  mtu 9216  
!  
interface xel.898  
  description for ### 6PE  
  encapsulation dot1q 898  
  load-interval 30  
  ipv6 address 3608::1/64  
  mtu 9216  
!  
interface xel.899  
  description for ### 6PE
```

```
encapsulation dot1q 899
load-interval 30
ipv6 address 3609::1/64
mtu 9216
!
interface xel.900
description for ### 6PE
encapsulation dot1q 900
load-interval 30
ipv6 address 360a::1/64
mtu 9216
!
interface xel.901 switchport
description for ### eline901
encapsulation dot1q 901
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772901
!
interface xel.902 switchport
description for ### eline902
encapsulation dot1q 902
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772902
!
interface xel.903 switchport
description for ### eline903
encapsulation dot1q 903
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772903
!
interface xel.904 switchport
description for ### eline904
encapsulation dot1q 904
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772904
!
interface xel.905 switchport
description for ### eline905
encapsulation dot1q 905
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772905
!
interface xel.906 switchport
description for ### eline906
encapsulation dot1q 906
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772906
!
interface xel.907 switchport
description for ### eline907
encapsulation dot1q 907
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 16772907
!
```

```
interface xe1.908 switchport
  description for ### eline908
  encapsulation dot1q 908
  load-interval 30
  mtu 9216
  access-if-evpn
    map vpn-id 16772908
!
interface xe1.909 switchport
  description for ### eline909
  encapsulation dot1q 909
  load-interval 30
  mtu 9216
  access-if-evpn
    map vpn-id 16772909
!
interface xe1.910 switchport
  description for ### eline910
  encapsulation dot1q 910
  load-interval 30
  mtu 9216
  access-if-evpn
    map vpn-id 16772910
!
interface xe1.1250 switchport
  description for ### bgp vpls1250
  encapsulation dot1q 1250
  load-interval 30
  mtu 9216
  access-if-vpls
    mpls-vpls vpls4294961250
!
interface xe1.1251 switchport
  description for ### bgp vpls1251
  encapsulation dot1q 1251
  load-interval 30
  mtu 9216
  access-if-vpls
    mpls-vpls vpls4294961251
!
interface xe1.1252 switchport
  description for ### bgp vpls1252
  encapsulation dot1q 1252
  load-interval 30
  mtu 9216
  access-if-vpls
    mpls-vpls vpls4294961252
!
interface xe1.1253 switchport
  description for ### bgp vpls1253
  encapsulation dot1q 1253
  load-interval 30
  mtu 9216
  access-if-vpls
    mpls-vpls vpls4294961253
!
interface xe1.1254 switchport
  description for ### bgp vpls1254
  encapsulation dot1q 1254
  load-interval 30
  mtu 9216
  access-if-vpls
    mpls-vpls vpls4294961254
!
interface xe1.1255 switchport
  description for ### bgp vpls1255
  encapsulation dot1q 1255
  load-interval 30
```

```
mtu 9216
access-if-vpls
  mpls-vpls vpls4294961255
!
interface xel.1256 switchport
description for ### bgp vpls1256
encapsulation dot1q 1256
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls4294961256
!
interface xel.1257 switchport
description for ### bgp vpls1257
encapsulation dot1q 1257
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls4294961257
!
interface xel.1258 switchport
description for ### bgp vpls1258
encapsulation dot1q 1258
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls4294961258
!
interface xel.1259 switchport
description for ### bgp vpls1259
encapsulation dot1q 1259
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls4294961259
!
interface xel.1260 switchport
description for ### bgp vpls1260
encapsulation dot1q 1260
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls4294961260
!
interface xel.1261 switchport
description for ### LDP_vpls1261
encapsulation dot1q 1261
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls1261
!
interface xel.1262 switchport
description for ### LDP_vpls1262
encapsulation dot1q 1262
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls1262
!
interface xel.1263 switchport
description for ### LDP_vpls1263
encapsulation dot1q 1263
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls vpls1263
!
```



```
interface xe1.1264 switchport
description for ### LDP_vpls1264
encapsulation dot1q 1264
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1264
!
interface xe1.1265 switchport
description for ### LDP_vpls1265
encapsulation dot1q 1265
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1265
!
interface xe1.1266 switchport
description for ### LDP_vpls1266
encapsulation dot1q 1266
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1266
!
interface xe1.1267 switchport
description for ### LDP_vpls1267
encapsulation dot1q 1267
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1267
!
interface xe1.1268 switchport
description for ### LDP_vpls1268
encapsulation dot1q 1268
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1268
!
interface xe1.1269 switchport
description for ### LDP_vpls1269
encapsulation dot1q 1269
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1269
!
interface xe1.1270 switchport
description for ### LDP_vpls1270
encapsulation dot1q 1270
load-interval 30
mtu 9216
access-if-vpls
mpls-vpls vpls1270
!
interface xe1.1271 switchport
description for ### mpls-l2-circuit-1271
encapsulation dot1q 1271
load-interval 30
mtu 9216
access-if-vpws
mpls-l2-circuit PE2-To-PE1-1271 primary
!
interface xe1.1272 switchport
description for ### mpls-l2-circuit-1272
encapsulation dot1q 1272
load-interval 30
```

```
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1272 primary
!
interface xel.1273 switchport
description for ### mpls-l2-circuit-1273
encapsulation dot1q 1273
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1273 primary
!
interface xel.1274 switchport
description for ### mpls-l2-circuit-1274
encapsulation dot1q 1274
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1274 primary
!
interface xel.1275 switchport
description for ### mpls-l2-circuit-1275
encapsulation dot1q 1275
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1275 primary
!
interface xel.1276 switchport
description for ### mpls-l2-circuit-1276
encapsulation dot1q 1276
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1276 primary
!
interface xel.1277 switchport
description for ### mpls-l2-circuit-1277
encapsulation dot1q 1277
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1277 primary
!
interface xel.1278 switchport
description for ### mpls-l2-circuit-1278
encapsulation dot1q 1278
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1278 primary
!
interface xel.1279 switchport
description for ### mpls-l2-circuit-1279
encapsulation dot1q 1279
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1279 primary
!
interface xel.1280 switchport
description for ### mpls-l2-circuit-1280
encapsulation dot1q 1280
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE1-1280 primary
!
```

```
interface xe1.1281 switchport
description for ### mpls-l2-circuit-1281
encapsulation dot1q 1281
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1281 primary
!
interface xe1.1282 switchport
description for ### mpls-l2-circuit-1282
encapsulation dot1q 1282
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1282 primary
!
interface xe1.1283 switchport
description for ### mpls-l2-circuit-1283
encapsulation dot1q 1283
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1283 primary
!
interface xe1.1284 switchport
description for ### mpls-l2-circuit-1284
encapsulation dot1q 1284
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1284 primary
!
interface xe1.1285 switchport
description for ### mpls-l2-circuit-1285
encapsulation dot1q 1285
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1285 primary
!
interface xe1.1286 switchport
description for ### mpls-l2-circuit-1286
encapsulation dot1q 1286
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1286 primary
!
interface xe1.1287 switchport
description for ### mpls-l2-circuit-1287
encapsulation dot1q 1287
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1287 primary
!
interface xe1.1288 switchport
description for ### mpls-l2-circuit-1288
encapsulation dot1q 1288
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1288 primary
!
interface xe1.1289 switchport
description for ### mpls-l2-circuit-1289
encapsulation dot1q 1289
load-interval 30
```

```
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1289 primary
!
interface xel.1290 switchport
description for ### mpls-l2-circuit-1290
encapsulation dot1q 1290
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE3-1290 primary
!
interface xel.1291 switchport
description for ### mpls-l2-circuit-1291
encapsulation dot1q 1291
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1291 primary
!
interface xel.1292 switchport
description for ### mpls-l2-circuit-1292
encapsulation dot1q 1292
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1292 primary
!
interface xel.1293 switchport
description for ### mpls-l2-circuit-1293
encapsulation dot1q 1293
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1293 primary
!
interface xel.1294 switchport
description for ### mpls-l2-circuit-1294
encapsulation dot1q 1294
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1294 primary
!
interface xel.1295 switchport
description for ### mpls-l2-circuit-1295
encapsulation dot1q 1295
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1295 primary
!
interface xel.1296 switchport
description for ### mpls-l2-circuit-1296
encapsulation dot1q 1296
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1296 primary
!
interface xel.1297 switchport
description for ### mpls-l2-circuit-1297
encapsulation dot1q 1297
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1297 primary
!
```

```
interface xe1.1298 switchport
description for ### mpls-l2-circuit-1298
encapsulation dot1q 1298
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1298 primary
!
interface xe1.1299 switchport
description for ### mpls-l2-circuit-1299
encapsulation dot1q 1299
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1299 primary
!
interface xe1.1300 switchport
description for ### mpls-l2-circuit-1300
encapsulation dot1q 1300
load-interval 30
mtu 9216
access-if-vpws
  mpls-l2-circuit PE2-To-PE4-1300 primary
!
interface xe1.1301 switchport
description for ### bgp-ad-1301
encapsulation dot1q 1301
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1301
!
interface xe1.1302 switchport
description for ### bgp-ad-1302
encapsulation dot1q 1302
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1302
!
interface xe1.1303 switchport
description for ### bgp-ad-1303
encapsulation dot1q 1303
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1303
!
interface xe1.1304 switchport
description for ### bgp-ad-1304
encapsulation dot1q 1304
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1304
!
interface xe1.1305 switchport
description for ### bgp-ad-1305
encapsulation dot1q 1305
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1305
!
interface xe1.1306 switchport
description for ### bgp-ad-1306
encapsulation dot1q 1306
load-interval 30
```

```
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1306
!
interface xel.1307 switchport
description for ### bgp-ad-1307
encapsulation dot1q 1307
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1307
!
interface xel.1308 switchport
description for ### bgp-ad-1308
encapsulation dot1q 1308
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1308
!
interface xel.1309 switchport
description for ### bgp-ad-1309
encapsulation dot1q 1309
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1309
!
interface xel.1310 switchport
description for ### bgp-ad-1310
encapsulation dot1q 1310
load-interval 30
mtu 9216
access-if-vpls
  mpls-vpls VPLS_AD1310
!
interface xel.1901 switchport
description for ### elan1901
encapsulation dot1q 1901
load-interval 30
mtu 9216
access-if-evpn
  map vpn-id 1901
!
interface xel.1902 switchport
description for ### elan1902
encapsulation dot1q 1902
load-interval 30
mtu 9216
access-if-evpn
  map vpn-id 61902
!
interface xel.1903 switchport
description for ### elan1903
encapsulation dot1q 1903
load-interval 30
mtu 9216
access-if-evpn
  map vpn-id 16777215
!
interface xel.1904 switchport
description for ### elan1904
encapsulation dot1q 1904
load-interval 30
mtu 9216
access-if-evpn
  map vpn-id 61904
!
```

```
interface xe1.1905 switchport
description for ### elan1905
encapsulation dot1q 1905
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 1671905
!
interface xe1.1906 switchport
description for ### elan1906
encapsulation dot1q 1906
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 1671906
!
interface xe1.1907 switchport
description for ### elan1907
encapsulation dot1q 1907
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 1671907
!
interface xe1.1908 switchport
description for ### elan1908
encapsulation dot1q 1908
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 1671908
!
interface xe1.1909 switchport
description for ### elan1909
encapsulation dot1q 1909
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 1671909
!
interface xe1.1910 switchport
description for ### elan1910
encapsulation dot1q 1910
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 1671910
!
interface xe1.3001 switchport
description for ### irbvrf3001
encapsulation dot1q 3001
rewrite pop
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 163001
!
interface xe1.3002 switchport
description for ### irbvrf3002
encapsulation dot1q 3002
rewrite pop
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 163002
!
interface xe1.3003 switchport
description for ### irbvrf3003
```

```
encapsulation dot1q 3003
rewrite pop
load-interval 30
mtu 9216
access-if-evpn
map vpn-id 163003
!
interface xe2
!
interface xe3
!
exit
!
router ospf 65535
fast-reroute keep-all-paths
shutdown
bfd all-interfaces
fast-reroute per-prefix remote-lfa area 0.0.0.1 tunnel mpls-ldp
network 192.168.36.2/32 area 0.0.0.1
network 203.0.113.10/31 area 0.0.0.1
network 203.0.113.16/31 area 0.0.0.1
network 203.0.113.18/31 area 0.0.0.1
!
router isis 1
is-type level-2-only
authentication mode md5 level-2
authentication key-chain isis level-2
ignore-lsp-errors
lsp-gen-interval 5
max-lsp-lifetime 2000
spf-interval-exp level-2 50 2000
metric-style wide
microloop-avoidance level-2
microloop-avoidance max-fib 60 level-2
mpls traffic-eng router-id 192.168.36.2
mpls traffic-eng level-2
capability cspf
dynamic-hostname
fast-reroute terminate-hold-on interval 100000
fast-reroute per-prefix level-2 proto ipv4 all
fast-reroute per-prefix remote-lfa level-2 proto ipv4 tunnel mpls-ldp
fast-reroute ti-lfa level-2 proto ipv4
bfd all-interfaces
net 49.0000.0000.0002.00
passive-interface loopback1
isis segment-routing global block 16000 23999
segment-routing entropy-label
!
router bgp 4294967294
bgp router-id 192.168.36.2
bgp auto-policy-soft-reset enable
bgp log-neighbor-changes
no bgp inbound-route-filter
allocate-label all
neighbor PEER-BGPLU peer-group
neighbor PEER-BGPLU remote-as 2200000002
neighbor PEER-RR peer-group
neighbor PEER-RR remote-as 4294967294
neighbor PEER-RR tcp-mss 1440
neighbor PEER-RR update-source loopback1
neighbor PEER-RR authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor PEER-RR advertisement-interval 0
neighbor PEER-RR fall-over bfd multihop
neighbor 190.160.2.254 peer-group PEER-BGPLU
neighbor 192.168.36.11 remote-as 4294967294
neighbor 192.168.36.11 tcp-mss 1440
neighbor 192.168.36.11 update-source loopback1
neighbor 192.168.36.11 authentication-key
```



```
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor 192.168.36.11 advertisement-interval 0
neighbor 192.168.36.11 fall-over bfd multihop
neighbor 192.168.36.12 peer-group PEER-RR
neighbor 3601::2 remote-as 200
neighbor 3602::2 remote-as 200
neighbor 3603::2 remote-as 200
neighbor 3604::2 remote-as 200
neighbor 3605::2 remote-as 200
neighbor 3606::2 remote-as 200
neighbor 3607::2 remote-as 200
neighbor 3608::2 remote-as 200
neighbor 3609::2 remote-as 200
neighbor 360a::2 remote-as 200
!
address-family ipv4 unicast
redistribute connected route-map REDISTRIBUTE-CONNECTED-TO-BGP
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
!
address-family ipv4 labeled-unicast
neighbor PEER-BGPLU activate
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
!
address-family vpnv4 unicast
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
!
address-family rtfilter unicast
exit-address-family
!
address-family l2vpn vpls
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
!
address-family l2vpn evpn
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
!
address-family vpnv6 unicast
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
!
address-family ipv6 unicast
redistribute connected
neighbor 3601::2 activate
neighbor 3602::2 activate
neighbor 3603::2 activate
neighbor 3604::2 activate
neighbor 3605::2 activate
neighbor 3606::2 activate
neighbor 3607::2 activate
neighbor 3608::2 activate
neighbor 3609::2 activate
neighbor 360a::2 activate
exit-address-family
!
address-family ipv6 labeled-unicast
neighbor PEER-RR activate
neighbor 192.168.36.11 activate
exit-address-family
```

```
!  
address-family ipv4 vrf irbvrf3001  
redistribute connected  
neighbor 80.10.1.254 remote-as 65535  
neighbor 80.10.1.254 activate  
neighbor 80.10.1.254 authentication-key 0x503653bfffef7c928057183d8be815ab  
exit-address-family  
!  
address-family ipv4 vrf irbvrf3002  
redistribute connected  
neighbor 80.11.1.254 remote-as 3002  
neighbor 80.11.1.254 activate  
neighbor 80.11.1.254 authentication-key 0x503653bfffef7c928057183d8be815ab  
exit-address-family  
!  
address-family ipv4 vrf irbvrf3003  
redistribute connected  
neighbor 80.12.1.254 remote-as 3003  
neighbor 80.12.1.254 activate  
neighbor 80.12.1.254 authentication-key 0x503653bfffef7c928057183d8be815ab  
exit-address-family  
!  
address-family ipv4 vrf vrf101  
redistribute connected  
redistribute static  
neighbor CLIENTS-V4 peer-group  
neighbor CLIENTS-V4 remote-as 65535  
neighbor CLIENTS-V4 activate  
neighbor CLIENTS-V4 authentication-key 0x503653bfffef7c928057183d8be815ab  
neighbor CLIENTS-V4 ebgp-multihop 255  
neighbor 201.101.1.2 peer-group CLIENTS-V4  
exit-address-family  
!  
address-family ipv4 vrf vrf102  
redistribute connected  
neighbor 201.102.1.2 remote-as 65535  
neighbor 201.102.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf103  
redistribute connected  
neighbor 201.103.1.2 remote-as 65535  
neighbor 201.103.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf104  
redistribute connected  
neighbor 201.104.1.2 remote-as 65535  
neighbor 201.104.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf105  
redistribute connected  
neighbor 201.105.1.2 remote-as 65535  
neighbor 201.105.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf106  
redistribute connected  
neighbor 201.106.1.2 remote-as 65535  
neighbor 201.106.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf107  
redistribute connected  
neighbor 201.107.1.2 remote-as 65535  
neighbor 201.107.1.2 activate  
exit-address-family
```

```
!  
address-family ipv4 vrf vrf108  
redistribute connected  
neighbor 201.108.1.2 remote-as 65535  
neighbor 201.108.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf109  
redistribute connected  
neighbor 201.109.1.2 remote-as 65535  
neighbor 201.109.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf110  
redistribute connected  
neighbor 201.110.1.2 remote-as 65535  
neighbor 201.110.1.2 activate  
exit-address-family  
!  
address-family ipv4 vrf vrf111  
neighbor CLIENTS peer-group  
neighbor CLIENTS remote-as 65534  
neighbor CLIENTS activate  
neighbor CLIENTS ebgp-multihop 2  
neighbor 201.111.1.2 remote-as 65535  
neighbor 201.111.1.2 activate  
neighbor 201.112.1.2 peer-group CLIENTS  
neighbor 201.113.1.2 remote-as 65533  
neighbor 201.113.1.2 activate  
neighbor 201.114.1.2 remote-as 65532  
neighbor 201.114.1.2 activate  
neighbor 201.115.1.2 remote-as 65531  
neighbor 201.115.1.2 activate  
neighbor 201.116.1.2 remote-as 65530  
neighbor 201.116.1.2 activate  
neighbor 201.117.1.2 remote-as 65529  
neighbor 201.117.1.2 activate  
neighbor 201.118.1.2 remote-as 65528  
neighbor 201.118.1.2 activate  
neighbor 201.119.1.2 remote-as 65527  
neighbor 201.119.1.2 activate  
neighbor 201.120.1.2 remote-as 65526  
neighbor 201.120.1.2 activate  
exit-address-family  
!  
address-family ipv6 vrf irbvrf3001  
redistribute connected  
neighbor 80:10:1::254 remote-as 65535  
neighbor 80:10:1::254 activate  
neighbor 80:10:1::254 authentication-key 0x503653bfffef7c928057183d8be815ab  
exit-address-family  
!  
address-family ipv6 vrf irbvrf3002  
redistribute connected  
neighbor 80:11:1::254 remote-as 3002  
neighbor 80:11:1::254 activate  
neighbor 80:11:1::254 authentication-key 0x503653bfffef7c928057183d8be815ab  
exit-address-family  
!  
address-family ipv6 vrf irbvrf3003  
redistribute connected  
neighbor 80:12:1::254 remote-as 3003  
neighbor 80:12:1::254 activate  
neighbor 80:12:1::254 authentication-key 0x503653bfffef7c928057183d8be815ab  
exit-address-family  
!  
address-family ipv6 vrf vrf101  
redistribute connected
```

```
neighbor CLIENTS peer-group
neighbor CLIENTS remote-as 65535
neighbor CLIENTS activate
neighbor CLIENTS authentication-key 0x503653bfffef7c928057183d8be815ab
neighbor CLIENTS ebgp-multihop 255
neighbor 2001:bd8:101::2 peer-group CLIENTS
exit-address-family
!
address-family ipv6 vrf vrf102
redistribute connected
neighbor 2001:bd8:102::2 remote-as 65535
neighbor 2001:bd8:102::2 activate
exit-address-family
!
address-family ipv6 vrf vrf103
redistribute connected
neighbor 2001:bd8:103::2 remote-as 65535
neighbor 2001:bd8:103::2 activate
exit-address-family
!
address-family ipv6 vrf vrf104
redistribute connected
neighbor 2001:bd8:104::2 remote-as 65535
neighbor 2001:bd8:104::2 activate
exit-address-family
!
address-family ipv6 vrf vrf105
redistribute connected
neighbor 2001:bd8:105::2 remote-as 65535
neighbor 2001:bd8:105::2 activate
exit-address-family
!
address-family ipv6 vrf vrf106
redistribute connected
neighbor 2001:bd8:106::2 remote-as 65535
neighbor 2001:bd8:106::2 activate
exit-address-family
!
address-family ipv6 vrf vrf107
redistribute connected
neighbor 2001:bd8:107::2 remote-as 65535
neighbor 2001:bd8:107::2 activate
exit-address-family
!
address-family ipv6 vrf vrf108
redistribute connected
neighbor 2001:bd8:108::2 remote-as 65535
neighbor 2001:bd8:108::2 activate
exit-address-family
!
address-family ipv6 vrf vrf109
redistribute connected
neighbor 2001:bd8:109::2 remote-as 65535
neighbor 2001:bd8:109::2 activate
exit-address-family
!
address-family ipv6 vrf vrf110
redistribute connected
neighbor 2001:bd8:110::2 remote-as 65535
neighbor 2001:bd8:110::2 activate
exit-address-family
!
address-family ipv6 vrf vrf111
redistribute connected
neighbor 2001:bd8:111::254 remote-as 65535
neighbor 2001:bd8:111::254 activate
neighbor 2001:bd8:112::254 remote-as 65534
neighbor 2001:bd8:112::254 activate
```

```

neighbor 2001:bd8:113::254 remote-as 65533
neighbor 2001:bd8:113::254 activate
neighbor 2001:bd8:114::254 remote-as 65532
neighbor 2001:bd8:114::254 activate
neighbor 2001:bd8:115::254 remote-as 65531
neighbor 2001:bd8:115::254 activate
neighbor 2001:bd8:116::254 remote-as 65530
neighbor 2001:bd8:116::254 activate
neighbor 2001:bd8:117::254 remote-as 65529
neighbor 2001:bd8:117::254 activate
neighbor 2001:bd8:118::254 remote-as 65528
neighbor 2001:bd8:118::254 activate
neighbor 2001:bd8:119::254 remote-as 65527
neighbor 2001:bd8:119::254 activate
neighbor 2001:bd8:11a::254 remote-as 65526
neighbor 2001:bd8:11a::254 activate
exit-address-family
!
exit
!
rsvp-trunk PE2_1_to_RR_1 ipv4
reoptimize
primary fast-reroute protection facility
primary fast-reroute node-protection
update-type make-before-break
to 192.168.36.12
!
rsvp-trunk PE2_1_to_RR2_1 ipv4
reoptimize
primary fast-reroute protection one-to-one
primary fast-reroute node-protection
primary label-record
update-type make-before-break
to 192.168.36.11
!
ip route vrf vrf101 0.0.0.0/0 Null
!
ip community-list standard CUSTOMER_ROUTES_3015_IRBVRF3003_IPv4 permit 51185:1015
ip community-list standard CUSTOMER_ROUTES_3015_IRBVRF3003_IPv6 permit 51185:3015
ip community-list standard CUSTOMER_ROUTES_ELAN1901 permit 1901:1901
!
ip large-community-list standard CUSTOMER_LARGE_COMM_ELAN1901 permit 1901:1901:65000
ip large-community-list standard CUSTOMER_LARGE_COMM_IRBVRF3003_V4 permit 200:3003:65000
ip large-community-list standard CUSTOMER_LARGE_COMM_IRBVRF3003_V6 permit 200:112:65000
!
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3016_IRBVRF3003_IPv4 permit rt 51185:4444
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3016_IRBVRF3003_IPv6 permit rt 51185:3333
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVRF3003_IPv4 permit rt 51185:65534
ip extcommunity-list standard CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVRF3003_IPv6 permit rt 51185:65534
ip extcommunity-list standard CUSTOMER_EXT_COMM_ELAN1901 permit rt 1901:1901
!
ip as-path access-list ASPATH-IRBVRF3003-IN-V4 permit ^65030$
ip as-path access-list ASPATH-IRBVRF3003-IN-V6 permit ^65020$
!
line console 0
exec-timeout 0
!
!
end

!
PE2-7001#

```

RR1

```

RR1-7036#sh run
!

```

```
! Software version: EC_AS5912-54X-OcNOS-SP-MPLS-7.0.0.129-Alpha 10/06/2025 17:41:16
!
! Last configuration change at 16:51:57 UTC Tue Oct 07 2025 by root
!
!
service password-encryption
!
logging console 5
logging monitor 5
logging cli
logging logfile ts_issue07 7
logging level nsm 3
logging level rip 5
logging level ripng 5
logging level ospf 5
logging level ospf6 5
logging level isis 5
logging level hostp 3
logging level ldp 5
logging level rsvp 5
logging level mrrib 5
logging level pim 5
logging level auth 5
logging level mstp 5
logging level onm 5
logging level hsl 3
logging level oam 5
logging level vlog 5
logging level vrrp 5
logging level ndd 5
logging level rib 5
logging level bgp 4
logging level l2mrrib 5
logging level lag 5
logging level sflow 5
logging level cml 3
logging level pserv 5
logging level cmm 4
logging level all 4
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
snmp-server enable traps pwdelete
snmp-server enable traps pw
snmp-server enable traps mpls
snmp-server enable traps snmp authentication
snmp-server enable traps ospf
snmp-server enable traps bgp
snmp-server enable traps isis
!
bgp extended-asn-cap
!
forwarding profile kaps profile-two
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
key chain isis
  key-id 3
  key-string encrypted 0xc8a471564ac751dc
!
key chain BGP
  key-id 4
  key-string encrypted 0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
!
qos enable
!
```

```
mpls lsp-stitching
mpls ilm-ecmp ldp
mpls ftn-ecmp ldp
mpls label mode vpnv4 all-vrfs per-vrf
mpls label mode vpnv6 all-vrfs per-vrf
mpls label mode all-afs all-vrfs per-vrf
!
hostname RR1-7036
ip domain-lookup vrf management
ip name-server vrf management 10.12.3.23
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
ospf restart grace-period 2
ospf restart helper max-grace-period 2
feature ssh vrf management
aaa local authentication attempts max-fail 25
aaa local authentication unlock-timeout 1
aaa authentication login error-enable
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community public vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
feature rsyslog
logging remote server 10.14.103.230 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514 vrf management
logging remote server 10.16.100.20 5 port 1514
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
router-id 12.12.12.12
!
evpn mpls enable
!
ip vrf management
!
segment-routing
!
ip multicast-routing
!
ipv6 multicast-routing
!
ip prefix-list PL-BGPLU
  seq 5 permit 101.101.101.101/32
  seq 10 permit 201.201.201.201/32
  seq 15 permit 13.13.13.13/32
  seq 20 permit 10.137.76.17/32
!
ip prefix-list PL-EVPN
  seq 5 permit 1.0.2.0/24
  seq 10 permit 1.0.1.0/24
!
ipv6 prefix-list PFX-EVPNV6
  seq 5 permit 2000:1:1:1::/64
  seq 10 permit 2000:1:1:2::/64
!
router ldp
  router-id 12.12.12.12
  fast-reroute
```

```
pw-status-tlv
!
router rsvp
!
route-map RM-EXPORT-BGPLU permit 10
  match ip address prefix-list PL-BGPLU
  set metric 333
  set community 1:1
!
route-map RM-EXPORT-BGPLU permit 20
!
route-map RM-EXPORT-EVPN permit 10
  match ip address prefix-list PL-EVPN
  set originator-id 11.12.13.14
!
route-map RM-EXPORT-EVPN permit 20
  match ipv6 address prefix-list PFX-EVPNv6
  set aigp-metric 234
  set originator-id 33.33.33.33
!
route-map RM-EXPORT-EVPN permit 50
!
interface ce49
  description connected_to_pe2
  load-interval 30
  ip address 203.0.113.18/31
  mtu 9194
  label-switching
  link-debounce-time 2000 0
  ip ospf network point-to-point
  ip ospf authentication message-digest
  ip ospf authentication-key 0xff87e79fdacd4e7
  ip ospf message-digest-key 3 md5 0x4c945d5d950eb831
  ipv6 ospf network point-to-point instance-id 0
  ipv6 router ospf area 0.0.0.0 tag 100 instance-id 0
  isis network point-to-point
  ip router isis 1
  ipv6 router isis 1
  isis authentication mode md5 level-1
  isis authentication mode md5 level-2
  isis authentication key-chain isis level-1
  isis authentication key-chain isis level-2
  enable-ldp ipv4
  mpls ldp-igp sync-delay 30
  enable-rsvp
  ip pim sparse-mode
  lldp-agent
    set lldp enable txrx
    set lldp chassis-id-tlv ip-address
    set lldp port-id-tlv if-name
    lldp tlv basic-mgmt system-name select
    lldp tlv basic-mgmt system-description select
  exit
  bfd interval 10 minrx 10 multiplier 3
!
interface ce50
!
interface ce51
!
interface ce52
!
interface ce53
!
interface ce54
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
```



```
!  
interface lo  
    ip address 127.0.0.1/8  
    ipv6 address ::1/128  
!  
interface lo.management  
    ip vrf forwarding management  
    ip address 127.0.0.1/8  
    ipv6 address ::1/128  
!  
interface loopback1  
    ip address 12.12.12.12/32  
    ipv6 address cafe:2012:12::12/128  
    prefix-sid index 7 explicit-null n-flag-clear  
    ipv6 router ospf area 0.0.0.0 instance-id 0  
    ip router isis 1  
    ipv6 router isis 1  
    ip pim sparse-mode  
!  
interface xe1  
!  
interface xe2  
!  
interface xe3  
!  
interface xe4  
!  
interface xe5  
!  
interface xe6  
!  
interface xe7  
!  
interface xe8  
!  
interface xe9  
!  
interface xe10  
!  
interface xe11  
!  
interface xe12  
!  
interface xe13  
!  
interface xe14  
!  
interface xe15  
!  
interface xe16  
!  
interface xe17  
!  
interface xe18  
!  
interface xe19  
!  
interface xe20  
!  
interface xe21  
!  
interface xe22  
!  
interface xe23  
!  
interface xe24  
!  
interface xe25
```

```
!  
interface xe26  
    description connected_to_pe1  
    load-interval 30  
    ip address 203.0.113.17/31  
    ipv6 address 2003:0:113::17/64  
    mtu 9194  
    label-switching  
    link-debounce-time 2000 0  
    isis network point-to-point  
    ip router isis 1  
    ipv6 router isis 1  
    isis authentication mode md5 level-1  
    isis authentication mode md5 level-2  
    isis authentication key-chain isis level-1  
    isis authentication key-chain isis level-2  
    enable-ldp ipv4  
    mpls ldp-igp sync-delay 30  
    enable-rsvp  
    ip pim sparse-mode  
    lldp-agent  
        set lldp enable txrx  
        set lldp chassis-id-tlv ip-address  
        set lldp port-id-tlv if-name  
        lldp tlv basic-mgmt system-name select  
        lldp tlv basic-mgmt system-description select  
    exit  
    bfd interval 10 minrx 10 multiplier 3  
!  
interface xe27  
!  
interface xe28  
!  
interface xe29  
!  
interface xe30  
!  
interface xe31  
!  
interface xe32  
!  
interface xe33  
!  
interface xe34  
!  
interface xe35  
!  
interface xe36  
!  
interface xe37  
!  
interface xe38  
!  
interface xe39  
!  
interface xe40  
!  
interface xe41  
!  
interface xe42  
!  
interface xe43  
!  
interface xe44  
!  
interface xe45  
!  
interface xe46
```

```
!  
interface xe47  
!  
interface xe48  
  description connected_to_cisco_port-5  
  load-interval 30  
  ip address 203.0.113.21/31  
  label-switching  
  isis network point-to-point  
  ip router isis 1  
  enable-ldp ipv4  
  lldp-agent  
    set lldp enable txrx  
    set lldp chassis-id-tlv ip-address  
    set lldp port-id-tlv if-name  
    lldp tlv basic-mgmt system-name select  
    lldp tlv basic-mgmt system-description select  
  exit  
  bfd interval 10 minrx 10 multiplier 3  
!  
  exit  
!  
router ospf 100  
  fast-reroute keep-all-paths  
  bfd all-interfaces  
  fast-reroute per-prefix remote-lfa area 0.0.0.0 tunnel mpls-ldp  
  network 12.12.12.12/32 area 0.0.0.0  
  network 203.0.113.18/31 area 0.0.0.0  
!  
router isis 1  
  is-type level-1  
  ignore-lsp-errors  
  lsp-gen-interval 5  
  max-lsp-lifetime 2000  
  spf-interval-exp level-2 50 2000  
  metric-style wide  
  microloop-avoidance level-1  
  microloop-avoidance max-fib 60 level-1  
  mpls traffic-eng router-id 12.12.12.12  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  fast-reroute terminate-hold-on interval 100000  
  fast-reroute per-prefix level-2 proto ipv4 all  
  fast-reroute per-prefix remote-lfa level-2 proto ipv4 tunnel mpls-ldp  
  fast-reroute ti-lfa level-2 proto ipv4  
  bfd all-interfaces  
  net 49.0001.0000.1102.00  
  isis segment-routing global block 16000 23999  
  segment-routing entropy-label  
!  
router isis ISIS-IGP-100  
  is-type level-1  
  authentication mode md5 level-1  
  authentication key-chain isis level-1  
  ignore-lsp-errors  
  lsp-gen-interval 5  
  max-lsp-lifetime 2000  
  spf-interval-exp level-1 50 2000  
  metric-style wide  
  microloop-avoidance level-1  
  microloop-avoidance max-fib 60 level-1  
  mpls traffic-eng router-id 12.12.12.12  
  mpls traffic-eng level-1  
  capability cspf  
  dynamic-hostname  
  fast-reroute terminate-hold-on interval 100000  
  fast-reroute per-prefix level-1 proto ipv4 all
```

```

fast-reroute ti-lfa level-1 proto ipv4
bfd all-interfaces
net 49.0001.0100.0000.1018.00
passive-interface
!
router isis OCNOS-CISCO
is-type level-1
ignore-lsp-errors
lsp-gen-interval 5
max-lsp-lifetime 2000
spf-interval-exp level-2 50 2000
metric-style wide
microloop-avoidance level-1
microloop-avoidance max-fib 60 level-1
mpls traffic-eng router-id 12.12.12.12
mpls traffic-eng level-1
capability cspf
dynamic-hostname
fast-reroute terminate-hold-on interval 100000
fast-reroute per-prefix level-2 proto ipv4 all
fast-reroute per-prefix remote-lfa level-2 proto ipv4 tunnel mpls-ldp
fast-reroute ti-lfa level-2 proto ipv4
bfd all-interfaces
net 49.0001.0000.0026.00
!
router bgp 4200000001
bgp router-id 12.12.12.12
bgp auto-policy-soft-reset enable
bgp cluster-id 4200000001
bgp log-neighbor-changes
no bgp inbound-route-filter
allocate-label all
neighbor PG-RR-PE1 peer-group
neighbor PG-RR-PE1 remote-as 4200000001
neighbor PG-RR-PE1 tcp-mss 1440
neighbor PG-RR-PE1 update-source loopback1
neighbor PG-RR-PE1 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor PG-RR-PE1 advertisement-interval 0
neighbor PG-RR-PE1 fall-over bfd multihop
neighbor PG-RR-PE2 peer-group
neighbor PG-RR-PE2 remote-as 4200000001
neighbor PG-RR-PE2 tcp-mss 1440
neighbor PG-RR-PE2 update-source loopback1
neighbor PG-RR-PE2 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor PG-RR-PE2 advertisement-interval 0
neighbor PG-RR-PE2 fall-over bfd multihop
neighbor 13.13.13.13 remote-as 65002
neighbor 13.13.13.13 ebgp-multihop 255
neighbor 13.13.13.13 update-source loopback1
neighbor 101.101.101.101 peer-group PG-RR-PE1
neighbor 201.201.201.201 peer-group PG-RR-PE2
!
address-family ipv4 unicast
network 12.12.0.12/32
network 12.12.12.12/32
exit-address-family
!
address-family ipv4 labeled-unicast
neighbor PG-RR-PE1 activate
neighbor PG-RR-PE1 route-reflector-client
neighbor PG-RR-PE1 route-map RM-EXPORT-BGPLU out
neighbor PG-RR-PE2 activate
neighbor PG-RR-PE2 route-reflector-client
neighbor PG-RR-PE2 route-map RM-EXPORT-BGPLU out
neighbor 13.13.13.13 activate
exit-address-family

```

```

!
address-family vpnv4 unicast
neighbor PG-RR-PE1 activate
neighbor PG-RR-PE1 route-reflector-client
neighbor PG-RR-PE1 next-hop-self
neighbor PG-RR-PE2 activate
neighbor PG-RR-PE2 route-reflector-client
neighbor PG-RR-PE2 next-hop-self
neighbor PG-RR-PE2 route-map RM-EXPORT-EVPN out
neighbor 13.13.13.13 allow-ebgp-vpn
neighbor 13.13.13.13 activate
neighbor 13.13.13.13 aigp enable
neighbor 13.13.13.13 route-map RM-EXPORT-EVPN out
exit-address-family
!
address-family rtfilter unicast
exit-address-family
!
address-family l2vpn vpls
neighbor PG-RR-PE1 activate
neighbor PG-RR-PE1 route-reflector-client
neighbor PG-RR-PE2 activate
neighbor PG-RR-PE2 route-reflector-client
exit-address-family
!
address-family l2vpn evpn
neighbor PG-RR-PE1 activate
neighbor PG-RR-PE1 route-reflector-client
neighbor PG-RR-PE2 activate
neighbor PG-RR-PE2 route-reflector-client
exit-address-family
!
address-family vpnv6 unicast
neighbor PG-RR-PE1 activate
neighbor PG-RR-PE1 route-reflector-client
neighbor PG-RR-PE2 activate
neighbor PG-RR-PE2 route-reflector-client
neighbor PG-RR-PE2 route-map RM-EXPORT-EVPN out
exit-address-family
!
address-family ipv6 unicast
redistribute connected
exit-address-family
!
address-family ipv6 labeled-unicast
neighbor PG-RR-PE1 activate
neighbor PG-RR-PE1 route-reflector-client
neighbor PG-RR-PE2 activate
neighbor PG-RR-PE2 route-reflector-client
exit-address-family
!
exit
!
!
end

!
RR1-7036#

```

RR1

```

RR#sh run
!
! Software version: EC_AS5912-54X-OcNOS-SP-MPLS-7.0.0.261-GA 02/20/2026 15:31:03
!
! Last configuration change at 17:14:02 UTC Tue Feb 24 2026 by root
!

```

```
service password-encryption
!
logging console 3
logging monitor 5
logging logfile device_debug_log 2
logging level nsm 5
logging level ospf 5
logging level ldp 5
logging level rsvp 5
logging level hsl 5
logging level bgp 5
logging level cml 5
logging level all 5
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
bgp extended-asn-cap
!
forwarding profile kaps profile-two
hardware-profile filter qos enable
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname RR
tfo Disable
errdisable cause stp-bpdu-guard
enable ext-ospf-multi-inst
feature dns relay
ip dns relay
ipv6 dns relay
!
ip vrf management
!
ip prefix-list PL-BGPLU
  seq 5 permit 1.1.1.1/32
  seq 10 permit 3.3.3.3/32
!
router ldp
  router-id 2.2.2.2
!
router rsvp
!
route-map RM-EXPORT-BGPLU permit 10
  match ip address prefix-list PL-BGPLU
  set ip next-hop self
!
route-map RM-EXPORT-BGPLU permit 20
!
interface ce49
  load-interval 30
  ip address 30.1.1.2/24
  mtu 9216
  label-switching
  ip ospf network point-to-point
  enable-rsvp
!
interface ce50
!
interface ce51
!
interface ce52
!
interface ce53
```

```
!  
interface ce54  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 2.2.2.2/32 secondary  
  ipv6 address ::1/128  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe1  
!  
interface xe2  
!  
interface xe3  
!  
interface xe4  
!  
interface xe5  
!  
interface xe6  
!  
interface xe7  
!  
interface xe8  
!  
interface xe9  
!  
interface xe10  
!  
interface xe11  
!  
interface xe12  
!  
interface xe13  
!  
interface xe14  
!  
interface xe15  
!  
interface xe16  
!  
interface xe17  
!  
interface xe18  
!  
interface xe19  
!  
interface xe20  
!  
interface xe21  
!  
interface xe22  
!  
interface xe23  
!  
interface xe24  
!  
interface xe25  
!  
interface xe26
```

```
load-interval 30
ip address 10.1.1.2/24
mtu 9216
label-switching
ip ospf network point-to-point
enable-ldp ipv4
!
interface xe27
!
interface xe28
!
interface xe29
!
interface xe30
!
interface xe31
!
interface xe32
!
interface xe33
!
interface xe34
!
interface xe35
!
interface xe36
!
interface xe37
!
interface xe38
!
interface xe39
!
interface xe40
!
interface xe41
!
interface xe42
!
interface xe43
!
interface xe44
!
interface xe45
!
interface xe46
!
interface xe47
!
interface xe48
!
exit
!
router ospf 65530
ospf router-id 2.2.2.2
bfd all-interfaces
network 2.2.2.2/32 area 0.0.0.0 instance-id 1
network 30.1.1.0/24 area 0.0.0.0
!
router ospf 65535
ospf router-id 2.2.2.2
bfd all-interfaces
network 2.2.2.2/32 area 0.0.0.0
network 10.1.1.0/24 area 0.0.0.0
!
router bgp 4200000001
bgp router-id 2.2.2.2
bgp auto-policy-soft-reset enable
```



```

bgp cluster-id 2.2.2.2
bgp log-neighbor-changes
no bgp inbound-route-filter
allocate-label all
neighbor 1.1.1.1 remote-as 4200000001
neighbor 1.1.1.1 tcp-mss 1440
neighbor 1.1.1.1 update-source 2.2.2.2
neighbor 1.1.1.1 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor 1.1.1.1 advertisement-interval 0
neighbor 1.1.1.1 fall-over bfd multihop
neighbor 3.3.3.3 remote-as 4200000001
neighbor 3.3.3.3 tcp-mss 1440
neighbor 3.3.3.3 update-source 2.2.2.2
neighbor 3.3.3.3 authentication-key
0xb8c718c634a41731bb38c629e7a365555c46a93e5e446157be7f6e35f32bf637
neighbor 3.3.3.3 advertisement-interval 0
neighbor 3.3.3.3 fall-over bfd multihop
!
address-family ipv4 unicast
network 2.2.2.2/32
exit-address-family
!
address-family ipv4 labeled-unicast
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 route-reflector-client
neighbor 1.1.1.1 route-map RM-EXPORT-BGPLU out
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 route-reflector-client
neighbor 3.3.3.3 route-map RM-EXPORT-BGPLU out
exit-address-family
!
address-family vpnv4 unicast
neighbor 1.1.1.1 activate
neighbor 1.1.1.1 route-reflector-client
neighbor 3.3.3.3 activate
neighbor 3.3.3.3 route-reflector-client
exit-address-family
!
exit
!
rsvp-trunk P1_to_PE3 ipv4
to 3.3.3.3
!
line console 0
exec-timeout 0 0
line vty 0 16
exec-timeout 0 0
!
!
end

```

Validation

Establish end-to-end PE1-PE2 labeled reachability using BGP-LU across dual OSPF domains (PE1-RR with OSPF+LDP and RR-PE2 with OSPF+RSVP)

Verification

```

## ! Define prefix-lists on RR node
RR#sh run prefix-list
!
ip prefix-list PL-BGPLU
seq 5 permit 1.1.1.1/32
seq 10 permit 3.3.3.3/32

```

```

!

## ! Define route-map for labeled-unicast

RR#sh run route-map
!
route-map RM-EXPORT-BGPLU permit 10
  match ip address prefix-list PL-BGPLU
  set ip next-hop self
!
route-map RM-EXPORT-BGPLU permit 20
!

# Verify prefix-list and route-map
RR#sh ip prefix-list detail PL-BGPLU
ip prefix-list PL-BGPLU:
  count: 2, range entries: 0, sequences: 5 - 10
  ripd:
    seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
    seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
  ripngd:
    seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
    seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
  ospfd:
    seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
    seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
  ospf6d:
    seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
    seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
  ldpd:
    seq 5 permit 1.1.1.1/32 (hit count: 0, refcount: 0)
    seq 10 permit 3.3.3.3/32 (hit count: 0, refcount: 0)
  bgpd:
    seq 5 permit 1.1.1.1/32 (hit count: 1, refcount: 1)
    seq 10 permit 3.3.3.3/32 (hit count: 1, refcount: 1)
RR#sh route
route-map router-id
RR#sh route-map
RR#sh route-map RM-EXPORT-BGPLU
route-map RM-EXPORT-BGPLU, permit, sequence 10
  Match clauses:
    ip address prefix-list: PL-BGPLU
  Set clauses:
    ip next-hop self
route-map RM-EXPORT-BGPLU, permit, sequence 20
  Match clauses:
  Set clauses:

## Verify on RR for next hop changed to itself (RR)
RR#sh ip bgp neighbors 1.1.1.1 advertised-routes

For address family: IPv4 Labeled-Unicast vrf: default
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

      Network          Next Hop          Metric    LocPrf    Weight Path    Ext-Color
*i>  2.2.2.2/32         2.2.2.2             0         100       32768 i         -

```

```

*>i  3.3.3.3/32      2.2.2.2      0      100      0      i      -
Total number of prefixes 2

RR#sh ip bgp labeled-unicast summary
BGP router identifier 2.2.2.2, local AS number 4200000001
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V      AS      MsgRcv      MsgSen  TblVer      InQ      OutQ      Up/Down  State/PfxRcd  Desc
1.1.1.1        4 4200000001      34        39         3         0         0  00:12:04      1
3.3.3.3        4 4200000001      35        43         3         0         0  00:12:33      1

Total number of neighbors 2

Total number of Established sessions 2

-----

RR#sh ip bgp neighbors 3.3.3.3 advertised-routes

For address family: IPv4 Labeled-Unicast vrf: default
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
                l - labeled, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
Description : Ext-Color - Extended community color

      Network      Next Hop      Metric      LocPrf      Weight Path      Ext-Color
*>i  1.1.1.1/32      2.2.2.2          0         100          0      i      -
*>i  2.2.2.2/32      2.2.2.2          0         100        32768      i      -
Total number of prefixes 2

RR#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled
Code  FEC      FTN-ID  Nhlf-ID  Tunnel-ID  Pri  Out-Label  Out-
Intf  ELC      Nexthop  Algo-Num  UpTime
L>  1.1.1.1/32      3        4        -        -        -        -
      -          N/A      00:19:32
      Yes  3          xe26      No      10.1.1.1      -        -
B  1.1.1.1/32      4        5        -        Yes  24324      -
      No      1.1.1.1      N/A      -
R
(t)> 3.3.3.3/32      1        1      5001      Yes  24320      ce49      No      3
0.1.1.1      N/A      00:20:01
B  3.3.3.3/32      2        2        -        Yes  24962      -
      No      3.3.3.3      N/A      -

RR#sh mpls ilm-table
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled
Code  FEC/VRF/L2CKT  ILM-ID  In-Label  Out-Label  In-Intf  Out-
Intf/VRF      Nexthop      pri  Algo-Num  UpTime  UPStr peers

```

```

B> 2.2.2.2/32      2      26240      Nolabel      N/A      N/A      127.0.0.1
    Yes N/A      00:19:57
R> 2.2.2.2/32      1      24320      Nolabel      N/A      N/A      127.0.0.1
    Yes N/A      00:19:59      1
L> 3.3.3.3/32      4      24961      Nolabel      N/A      N/A      127.0.0.1
    Yes N/A      00:19:36      1
B> 3.3.3.3/32      3      26241      24962      N/A      N/A      3.3.3.3
    Yes N/A      00:19:54
B> 1.1.1.1/32      5      26242      24324      N/A      N/A      1.1.1.1
    Yes N/A      00:19:25
RR#

=====

## Verification on PE1 Next hop changed to RR Loopback address

PE1#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal, S - stale
      Network      Next Hop      In Label      Out Label
*> 1.1.1.1/32      0.0.0.0      24324      -
*>i 2.2.2.2/32      2.2.2.2      24323      26240
*>i 3.3.3.3/32      2.2.2.2      24322      26241
PE1#sh ip bgp labeled-unicast summary
BGP router identifier 1.1.1.1, local AS number 4200000001
BGP table version is 3
1 BGP AS-PATH entries
0 BGP community entries

Neighbor      V      AS      MsgRcv      MsgSen      TblVer      InQ      OutQ      Up/Down      State/PfxRcd      Desc
2.2.2.2      4      4200000001      43      38      2      0      0 00:13:41      2

Total number of neighbors 1

Total number of Established sessions 1

PE1#sh ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
      O - OSPF, IA - OSPF inter area
      N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
      E1 - OSPF external type 1, E2 - OSPF external type 2
      i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
      ia - IS-IS inter area, E - EVPN,
      v - vrf leaked
      * - candidate default

IP Route Table for VRF "default"
C      1.1.1.1/32 is directly connected, lo, installed 00:14:39, last update 00:14:39 ago
O      2.2.2.2/32 [110/2] via 10.1.1.2, xe26, installed 00:14:25, last update 00:14:25 ago
B      3.3.3.3/32 [200/0] via 2.2.2.2 (recursive via 10.1.1.2), installed 00:04:33, last update
00:04:33 ago
C      10.1.1.0/24 is directly connected, xe26, installed 00:14:40, last update 00:14:40 ago
C      127.0.0.0/8 is directly connected, lo, installed 00:22:06, last update 00:22:06 ago

Gateway of last resort is not set

PE1#sh mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
      B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
      L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
      U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
      (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Disabled

Code	FEC	FTN-ID	Nhlfe-ID	Tunnel-ID	Pri	Out-Label	Out-
Intf	ELC	Nexthop	Algo-Num	UpTime			
L>	2.2.2.2/32	3	6	-	-	-	-
	-	N/A	00:20:03				
	Yes	3	xe26	No		10.1.1.2	-
B	2.2.2.2/32	2	4	-	Yes	26240	-
	No	2.2.2.2	N/A	-			
B>	3.3.3.3/32	1	12	-			
	Yes	26241	xe26	No		2.2.2.2	N/A
							00:10:31

PE1#sh mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
 K - CLI ILM, T - MPLS-TP, s - Stitched ILM
 S - SNMP, L - LDP, R - RSVP, C - CRLDP
 B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
 O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
 P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled

Code	FEC/VRF/L2CKT	ILM-ID	In-Label	Out-Label	In-Intf	Out-
Intf/VRF	Nexthop		pri	Algo-Num	UpTime	UPStr peers
B>	VRF2	2	24321	Nolabel	N/A	N/A
	Yes N/A	00:20:13				N/A
B>	VRF1	1	24320	Nolabel	N/A	N/A
	Yes N/A	00:20:13				N/A
B>	2.2.2.2/32	4	24323	26240	N/A	N/A
	Yes N/A	00:20:06				2.2.2.2
B>	3.3.3.3/32	3	24322	26241	N/A	N/A
	Yes N/A	00:10:34				2.2.2.2
B>	1.1.1.1/32	5	24324	Nolabel	N/A	N/A
	Yes N/A	00:20:09				127.0.0.1

PE1#ping ip 3.3.3.3

Press CTRL+C to exit

PING 3.3.3.3 (3.3.3.3) 100(128) bytes of data.

108 bytes from 3.3.3.3: icmp_seq=1 ttl=64 time=0.641 ms

108 bytes from 3.3.3.3: icmp_seq=2 ttl=64 time=0.537 ms

108 bytes from 3.3.3.3: icmp_seq=3 ttl=64 time=0.669 ms

108 bytes from 3.3.3.3: icmp_seq=4 ttl=64 time=0.546 ms

108 bytes from 3.3.3.3: icmp_seq=5 ttl=64 time=0.606 ms

--- 3.3.3.3 ping statistics ---

5 packets transmitted, 5 received, 0% packet loss, time 4132ms

rtt min/avg/max/mdev = 0.537/0.599/0.669/0.051 ms

Verify on PE2

PE2#sh ip bgp labeled-unicast

Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i - internal, S - stale

Network	Next Hop	In Label	Out Label
*>i 1.1.1.1/32	2.2.2.2	24964	26242
*>i 2.2.2.2/32	2.2.2.2	24963	26240
*> 3.3.3.3/32	0.0.0.0	24962	-

PE2#sh ip bgp labeled-unicast summary

BGP router identifier 3.3.3.3, local AS number 4200000001

BGP table version is 4

1 BGP AS-PATH entries

0 BGP community entries

Neighbor	V	AS	MsgRcv	MsgSen	TblVer	InQ	OutQ	Up/Down	State/PfxRcd	Desc
2.2.2.2	4	4200000001		47	40	3	0	0 00:14:38		2

Total number of neighbors 1

Total number of Established sessions 1

PE2#sh ip route

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,

ia - IS-IS inter area, E - EVPN,

v - vrf leaked

* - candidate default

IP Route Table for VRF "default"

B 1.1.1.1/32 [200/0] via 2.2.2.2 (recursive via 30.1.1.2), installed 00:03:59, last update 00:03:59 ago

O 2.2.2.2/32 [110/2] via 30.1.1.2, ce0, installed 00:14:44, last update 00:14:44 ago

C 3.3.3.3/32 is directly connected, lo, installed 00:14:59, last update 00:14:59 ago

C 30.1.1.0/24 is directly connected, ce0, installed 00:14:59, last update 00:14:59 ago

C 127.0.0.0/8 is directly connected, lo, installed 00:21:32, last update 00:21:32 ago

Gateway of last resort is not set

PE2#sh mpls forwarding-table

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup

B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,

L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,

U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN

(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Disabled

Code	FEC	Nexthop	FTN-ID	Nhlfe-ID	Tunnel-ID	Pri	Out-Label	Out-
Intf	ELC		Algo-Num	UpTime				
B>	1.1.1.1/32	3	11	-	2.2.2.2		N/A	00:11:11
	Yes	26242	ce0	No				
R								
(t)>	2.2.2.2/32	1	3	5001	Yes	24320	ce0	No
	0.1.1.2	N/A	00:21:09					3
B	2.2.2.2/32	2	4	-	Yes	26240	-	
	No	2.2.2.2	N/A	-				

PE2#sh mpls ilm-table

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup

K - CLI ILM, T - MPLS-TP, s - Stitched ILM

S - SNMP, L - LDP, R - RSVP, C - CRLDP

B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT

O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI

P - SR Policy, U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Disabled

Code	FEC/VRF/L2CKT	ILM-ID	In-Label	Out-Label	In-Intf	Out-
Intf/VRF	Nexthop		pri	Algo-Num	UpTime	UPStr peers
B>	3.3.3.3/32	4	24962	Nolabel	N/A	N/A
	Yes N/A	00:21:12				127.0.0.1
B>	VRF1	1	24960	Nolabel	N/A	N/A
	Yes N/A	00:21:22				N/A
R>	3.3.3.3/32	3	24320	Nolabel	N/A	N/A
	Yes N/A	00:21:18	1			127.0.0.1
B>	VRF2	2	24961	Nolabel	N/A	N/A
	Yes N/A	00:21:22				N/A
B>	2.2.2.2/32	5	24963	26240	N/A	N/A
	Yes N/A	00:21:06				2.2.2.2
B>	1.1.1.1/32	6	24964	26242	N/A	N/A
	Yes N/A	00:11:14				2.2.2.2

PE2#ping ip 1.1.1.1

Press CTRL+C to exit

PING 1.1.1.1 (1.1.1.1) 100(128) bytes of data.

108 bytes from 1.1.1.1: icmp_seq=1 ttl=64 time=0.672 ms

```
108 bytes from 1.1.1.1: icmp_seq=2 ttl=64 time=0.491 ms
108 bytes from 1.1.1.1: icmp_seq=3 ttl=64 time=0.516 ms
108 bytes from 1.1.1.1: icmp_seq=4 ttl=64 time=0.515 ms
108 bytes from 1.1.1.1: icmp_seq=5 ttl=64 time=0.516 ms

--- 1.1.1.1 ping statistics ---
5 packets transmitted, 5 received, 0% packet loss, time 4103ms
rtt min/avg/max/mdev = 0.491/0.542/0.672/0.065 ms
```

Commands

The BGP LU Next-hop self in Route-map feature has the following configuration commands:

set ip next-hop self

Use this command to set next hop self for IPV4 BGP-LU neighbors.

Use “no” form of this command to not set next hop self.

Command Syntax

```
set ip next-hop self
no set ip next-hop self
```

Parameters

None

Command Mode

Route map mode mode

Applicability

This command is introduced in OcNOS version 7.0.0

Examples

```
ip prefix-list BCOM-IP
  seq 5 permit 25.4.4.0/24 eq 24

route-map BCOM-RM permit 2
  match ip address prefix-list BCOM-IP
  set ip next-hop self

router bgp 26
..
  address-family ipv4 labeled-unicast
  neighbor 27.27.27.27 route-map BCOM-RM out
..
```

| LAYER 2 SERVICE ENHANCEMENTS

Enhancements in Layer 2 services improve link reliability, convergence, and operational control for Ethernet-based deployments. The updates optimize service resiliency and simplify network operations in access and aggregation layers.

EVPN MPLS E-Tree Scenario 2	177
Overview	177
E-Tree Scenario 2 Prerequisites	179
E-Tree Scenario 2 Configuration	182
E-Tree Scenario 2 Implementation Examples	207
E-Tree Scenario 2 Troubleshooting	207
E-Tree Scenario 2 Revised Commands	208
E-Tree Scenario 2 Glossary	210

EVPN MPLS E-Tree Scenario 2

Overview

Ethernet VPN Ethernet-Tree (EVPN E-Tree) Scenario 2 (SC-2) [RFC 8317] enables Root and Leaf sites to co-exist on the same Provider Edge (PE) device within a single EVPN Instance (EVI). The solution enforces traffic filtering rules both within a PE (intra-PE) and across PEs (inter-PE), ensuring strict Leaf-to-Leaf communication control while maintaining flexible service deployment.

Scenario 2 supports both single-homed (SH) and multi-homed (MH) configurations and applies to Qumran2 (Q2) series platforms. All other EVPN E-Tree functionality and route types conform to [RFC 7432], except where extended for Scenario 2 behavior.

Inter-PE E-tree: New Route Advertisements

Scenario 2 introduces the following new route advertisements to support Leaf traffic isolation across PEs:

- **Leaf Host Advertisement (RT-2):** RT-2 advertisements for Leaf hosts include a new E-Tree Extended Community attribute. This attribute identifies the host as a Leaf for unicast traffic handling.
- **EAD per ES Advertisement (RT-1):** A new EAD per ES Route Type-1 (RT-1) is advertised with ESI=0. This route carries the E-Tree Extended Community attribute, which contains the Leaf Label. This mechanism is used to filter unknown traffic between a leaf source and a leaf destination.

Feature Characteristics

EVPN E-Tree Scenario 2 allows Root and Leaf Attachment Circuits (ACs) to exist within the same EVI on a PE. Traffic between Leaf ACs is restricted both locally and across PEs, maintaining E-Tree hierarchy and service separation.

Intra-PE Traffic Filtering

OcNOS filters traffic locally between Leaf ACs that belong to the same EVI on a single PE.

Traffic Flow (Unicast or BUM):

Leaf AC to Leaf AC (same PE): Traffic originating from one Leaf AC and destined for another Leaf AC on the same PE is dropped. This is due to the enforcement of split-horizon, which prevents Leaf-to-Leaf communication within the same PE. Any unicast or unknown traffic is filtered based on the source and destination ACs; if both are Leaf ACs, the traffic is dropped.

Inter-PE Traffic Filtering (Leaf ACs in the Same EVI)

To restrict communication between Leaf Access Circuits (ACs) located on different Provider Edges (PEs) within the same EVPN Instance (EVI), the amended routes described previously are utilized.

Unicast Traffic Filtering

Traffic filtering for Unicast traffic between Leaf ACs is performed at the originating Leaf AC PE node (the ingress PE node). This is possible because the host advertised from a Leaf AC at a remote PE is identified as a Leaf through the new RT-2 E-TREE attribute. Consequently, the Local PE recognizes the remote Host as one advertised from a Leaf AC, enabling ingress filtering. Reference: [RFC 8317, Section 4.1](#).

BUM Traffic Filtering

- Traffic filtering for Unknown Broadcast, Unknown Unicast, Multicast (BUM) traffic between Leaf ACs is implemented at the egress PE device. In the case of BUM traffic, the receivers include all ACs in the EVPN instance, some of which may be Root and others Leaf. Therefore, filtering cannot be restricted at the ingress PE.
- The remote PE uses RT-1 with ESI=0 and the E-TREE attribute to advertise a Leaf Label. The local Leaf AC uses this Leaf Label when sending BUM traffic.

Reference: [RFC 8317, Section 4.2](#)

Handling Multi-Homing (MH)

In Multi-Homing scenarios, if a Leaf AC is also on an ESI Multi-Homing port, the Leaf Label is prioritized over the ESI Label for advertising.

Intra-PE Traffic Flow Details

- **Co-existing Root and Leaf Sites on the Same PE:** A given Provider Edge (PE) may simultaneously host both Root and Leaf Attachment Circuits (ACs) for a specific Ethernet Virtual Instance (EVI).
- **Ingress Filtering (Unicast or Known Traffic):**
 - Traffic is subjected to filtering upon entering the PE via a Leaf AC.
 - A dedicated grouping identifies all Media Access Control (MAC) addresses learned from remote Leaf ACs.
 - Leaf-to-Leaf unicast traffic is discarded to prevent unauthorized communication.
- **Egress Filtering (BUM Traffic):**
 - BUM traffic includes a Leaf Label, which is advertised by the remote PE.
 - The Egress PE utilizes the Leaf Label to inhibit traffic transmission towards Leaf ACs.
 - This mechanism is effective for both single-homed (SH) and multi-homed (MH) Leaf ACs.
- **Route Exchange and Attributes:**

Route Type	Attribute Function	Filtering Scope
RT-2 (MAC/IP Advertisement)	Carries Leaf indication for ingress unicast filtering	Ingress PE
RT-1 (Ethernet A-D per ES, ESI=0)	Carries Leaf Label for egress BUM filtering	Egress PE



Note: The Leaf Label is scoped per PE, not per EVI or per ES.

Traffic Filtering Rules for All Combinations

Unicast

- **Leaf AC to Leaf MAC:** Traffic originating from a Leaf AC and destined for a MAC address learned from any Leaf AC on a remote PE is dropped, thereby preventing Leaf-to-Leaf communication across PEs.
- **Root AC to Leaf MAC:** Traffic is permitted, facilitating Root-to-Leaf communication.
- **Root AC to non-Leaf (Root) MAC:** Traffic is permitted, facilitating Root-to-Root communication.
- **Leaf AC to non-Leaf (Root) MAC:** Traffic is permitted, facilitating Leaf-to-Root communication.

BUM

- **SH Leaf AC:** Traffic is tagged with the Leaf Label; the egress PE drops it toward Leaf ACs.
- **MH Leaf AC on ESI port:** Traffic carries the Leaf Label; the egress PE blocks delivery to all Leaf ACs (both SH and MH).
- **Leaf-to-Leaf traffic within the same PE:** Traffic is dropped (via split-horizon functionality).



Note: Configure `hardware-profile filter evpn-mpls-mh` group even for single-homing nodes to enable Leaf Label enforcement.

Benefits

- **Strict traffic enforcement:** Leaf-to-Leaf communication is always blocked, maintaining E-Tree hierarchy and isolation.
- **Flexible service deployment:** Root and Leaf sites can co-exist on the same PE, simplifying design and reducing the number of EVIs.
- **Simplified control-plane operation:** A single route type per EVI with the E-Tree Extended Community and Leaf Label reduces BGP processing complexity.
- **Enhanced traffic security:** Ingress and egress filtering ensure unauthorized or misrouted traffic is dropped at the hardware level.
- **Consistent SH or MH operation:** Unified filtering logic supports both single-homed and multi-homed Leaf ACs.
- **Efficient hardware utilization:** Leaf Label allocation per PE enables hardware-level enforcement without additional per-EVI or per-ES labels.

E-Tree Scenario 2 Prerequisites

Before configuring E-Tree Scenario-2 in [Figure 5](#) topology, ensure the following prerequisites:

Root PEs and Leaf PEs**Underlay Transport Setup (MPLS, LDP, RSVP, IGP)**

MPLS, LDP, RSVP, and IS-IS (or OSPF) are operational across all PE and core routers. These protocols establish the MPLS underlay for label distribution and traffic forwarding.

```
!
router ospf 100
  ospf router-id 1.1.1.1
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 12.1.1.0/24 area 0.0.0.0
  network 13.1.1.0/24 area 0.0.0.0
!
evpn mpls vtep-ip-global 1.1.1.1
!
evpn mpls id 100
  host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
```

```

    host-reachability-protocol evpn-bgp vrf101
    !
    evpn mpls id 102
    host-reachability-protocol evpn-bgp vrf102
    !
    router ldp
    graceful-restart full
    targeted-peer ipv4 2.2.2.2
    exit-targeted-peer-mode
    targeted-peer ipv4 3.3.3.3
    exit-targeted-peer-mode
    targeted-peer ipv4 4.4.4.4
    exit-targeted-peer-mode
    transport-address ipv4 1.1.1.1
    !
    router rsvp
    !
    rsvp-path PE1-PE3 mpls
    13.1.1.2 strict
    17.1.1.1 strict
    19.1.1.2 strict
    !
    rsvp-path PE1-PE4 mpls
    12.1.1.2 strict
    17.1.1.2 strict
    21.1.1.2 strict
    !
    rsvp-trunk PE1-PE3 ipv4
    primary fast-reroute protection one-to-one
    primary path PE1-PE3
    to 3.3.3.3
    !
    rsvp-trunk PE1-PE4 ipv4
    primary fast-reroute protection facility
    primary path PE1-PE4
    to 4.4.4.4
    !
    rsvp-trunk PE1-PE2 ipv4
    to 2.2.2.2
    !

```

Loopback Interface Reachability

Loopback interfaces on each PE and core router are reachable through the IGP. Loopbacks provide router IDs for BGP EVPN and LDP sessions.

```

!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
 ip router isis ISIS-IGP
 enable-rsvp
!

```

Global EVPN and Multihoming Enablement

EVPN MPLS and Multihoming are globally enabled and allow multi-homed Ethernet segments for redundancy.

```

!
evpn mpls enable
!
evpn esi hold-time 200
!
evpn mpls multihoming enable
!

```

BGP EVPN Control Plane Establishment

BGP EVPN sessions are established among all PE routers. Each PE's loopback interface is used for BGP EVPN update-source and RSVP destination.

```
!  
router bgp 65010  
 neighbor EVPN peer-group  
 neighbor EVPN remote-as 65010  
 neighbor EVPN update-source lo  
 neighbor EVPN advertisement-interval 0  
 neighbor EVPN fall-over bfd multihop  
 neighbor 2.2.2.2 peer-group EVPN  
 neighbor 3.3.3.3 peer-group EVPN  
 neighbor 4.4.4.4 peer-group EVPN  
!  
 address-family l2vpn evpn  
 neighbor EVPN activate  
 exit-address-family  
!  
 exit  
!
```

Core-Facing Interface Configuration

Port-channels or physical links toward the core have MPLS, LDP, RSVP, and IGP enabled. Each core-facing link supports transport signaling for MPLS services.

```
!  
interface po10  
 description connected-to-p1  
 load-interval 30  
 ip address 12.1.1.1/24  
 mtu 9216  
 label-switching  
 enable-ldp ipv4  
 enable-rsvp  
!  
interface po20  
 description connected-to-p2  
 load-interval 30  
 ip address 13.1.1.1/24  
 mtu 9216  
 label-switching  
 enable-ldp ipv4  
 enable-rsvp  
!
```

MPLS Core Fabric Formation (P1 or P2 Routers)

[Core Routers \(page 201\)](#) (P1 and P2) form the MPLS transport fabric between the root and leaf PEs (Core routers interconnect the EVPN PEs and carry labeled traffic.).

```
!  
interface xe1  
 description connected-to-p2  
 load-interval 30  
 ip address 17.1.1.1/24  
 mtu 9216  
 label-switching  
 enable-ldp ipv4
```

```
enable-rsvp
!
```

Access Switch VLAN Trunk Configuration

[Access Switches \(page 206\)](#) (SW-1 and SW-2) are configured with VLAN trunks toward PEs carrying tagged leaf traffic.

```
!
interface po1000
  switchport
  bridge-group 1 spanning-tree disable
  switchport mode trunk
  switchport trunk allowed vlan all
!
interface xe3/1
  port breakout enable 4X10g
  switchport
  bridge-group 1 spanning-tree disable
  switchport mode trunk
  switchport trunk allowed vlan all
!
```

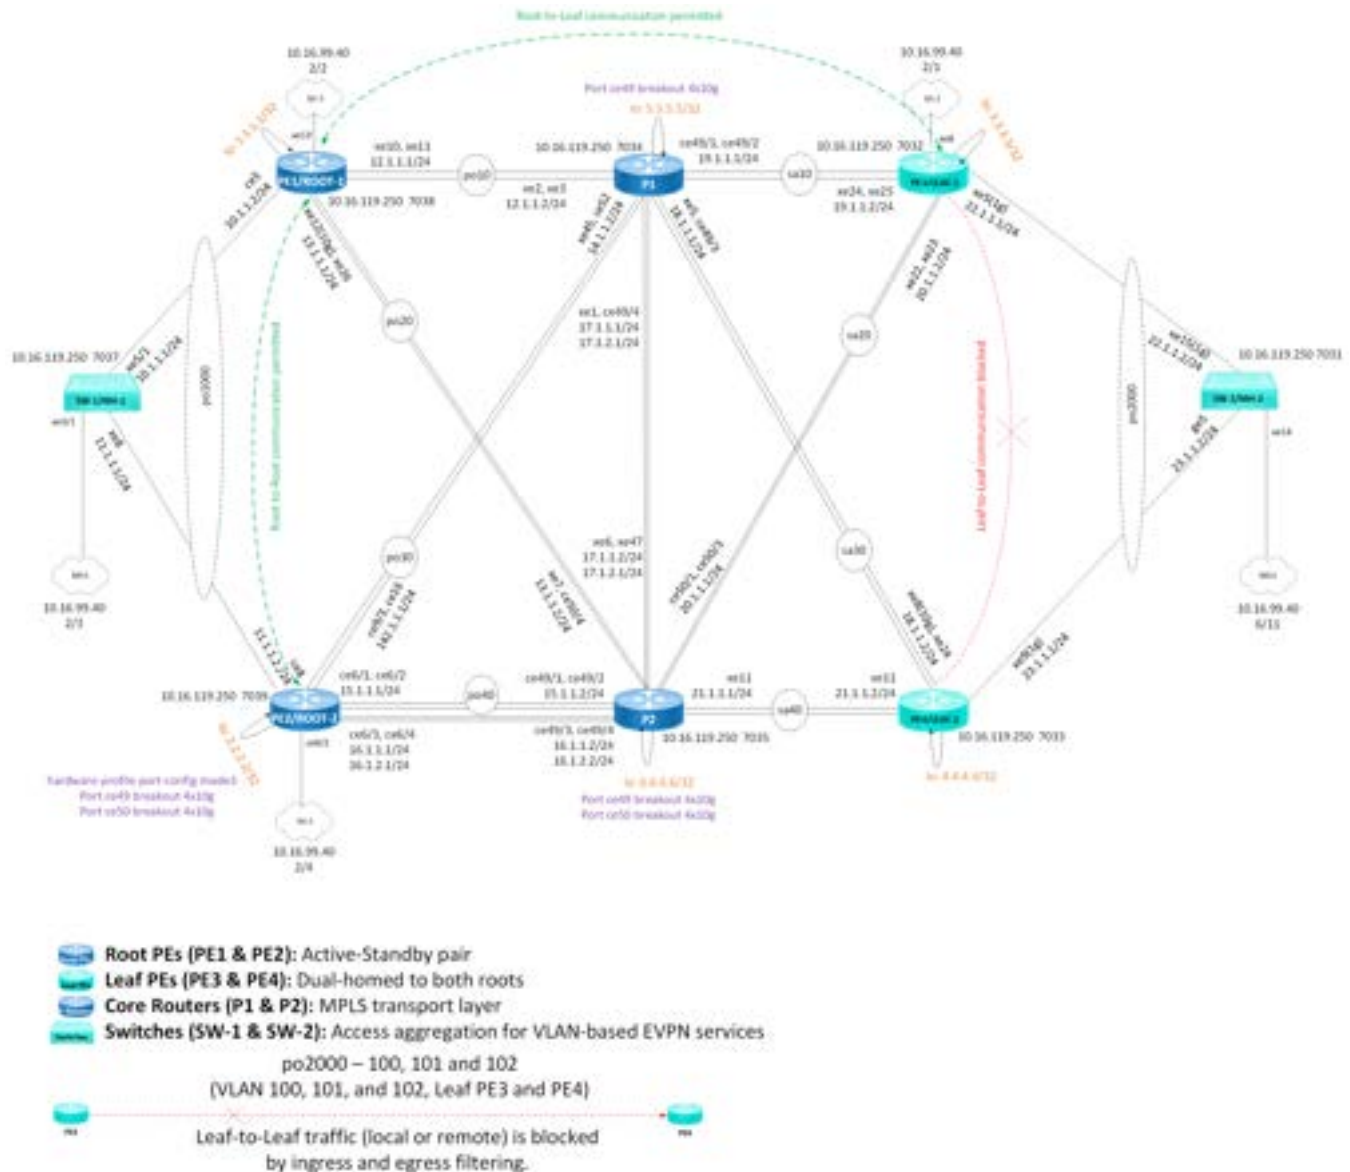
E-Tree Scenario 2 Configuration

This section describes a sample use case topology illustrating the E-Tree Scenario 2 deployment in an EVPN MPLS network. It explains enabling the scenario-2 on Provider Edge (PE) routers, its operation across the MPLS core, and validation steps to ensure proper functionality and traffic distribution.

Topology

Node	Role or Function in EVPN E-Tree SC-2
PE1 (Root)	Acts as a Root PE; originates and receives traffic from all Leaf and Root sites. Advertises RT-2 routes with Leaf indication for filtering.
PE2 (Root)	Redundant Root PE (active-standby pair with PE1). Provides Root-to-Root and Root-to-Leaf communication. Enforces ingress or egress filtering.
PE3 (Leaf)	Leaf PE (single-homed or multi-homed). Learns MACs from Leaf ACs and marks them with Leaf attributes. Cannot send directly to other Leaf ACs.
PE4 (Leaf)	Dual-homed Leaf PE (multi-homed with ESI). Sends BUM traffic with Leaf Label; remote PEs use label to drop traffic to Leaf sites.
P1 (Core)	MPLS core router transporting labeled EVPN traffic between PEs. Does not participate in E-Tree filtering but forwards based on labels.
P2 (Core)	MPLS core router providing redundancy and ECMP forwarding for EVPN or MPLS traffic.
SW-1 (Access)	Provides VLAN-based connectivity to Leaf subscribers; connects to Leaf ACs. Classified as Leaf-facing interface.
SW-2 (Access)	Provides VLAN-based connectivity to Root-facing services (e.g., data center). Classified as Root-facing interface.

Figure 5. E-Tree Scenario 2



Note: Before configuration meet all [E-Tree Scenario 2 Prerequisites](#) (page 179).

Enable E-Tree Scenario-2 Functionality

E-Tree Scenario-2 is enabled globally using the `evpn etree enable scenario-2` command. Activates E-Tree mode allowing both Root and Leaf sites on the same PE.

```
!
evpn etree enable scenario-2
!
hardware-profile filter evpn-mpls-mh enable
!
```

VPN ID and Leaf Role Mapping

- VPN IDs are defined for each E-Tree service, and access sub-interfaces are mapped with `map vpn-id <id> etree-leaf` where leaf identification is required. Associates sub-interfaces with E-Tree roles (Root or Leaf).
- Root sites communicate with all others; Leaf sites are isolated from each other. Configured using the `etree-leaf` parameter on access interfaces.

```
!
interface po1000
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.1212.1313
!
interface po1000.100 switchport
  encapsulation dot1q 100
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 100
!
interface po1000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 101 etree-leaf
!
interface po1000.102 switchport
  encapsulation dot1q 102
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 102 etree-leaf
!
```

E-Tree Scenario 2 Validation

After configuring the E-Tree Scenario 2 feature on all Provider Edge (PE) nodes, perform the following validation steps to confirm correct operation of ingress or egress filtering, Leaf Label advertisement, and BGP route exchange.



Notes:

- Repeat the same verification on all PEs (PE1–PE4) to confirm consistent Leaf Label allocation and DF or NON-DF synchronization.
- Ensure the `evpn etree enable scenario-2` and `map vpn-id <ID> etree-leaf` configurations are active before validation.

Verify DF and Non-DF Status on Multi-Homed PEs

- Displays per-EVI Designated Forwarder (DF) or Non-Designated Forwarder (Non-DF) status for multi-homed access circuits (ACs).
- In Scenario 2, DF or Non-DF roles ensure that only the DF PE forwards BUM traffic toward the shared Ethernet Segment Identifier (ESI), preventing duplication.
- Each VLAN or sub-interface under the same ESI should display alternating **DF** or **NON-DF** roles between MH-PEs.
- The below output confirms deterministic BUM handling across redundant PEs.


```

PE1#show evpn mpls
EVPN-MPLS Information
=====
Codes: NW - Network Port
      AC - Access Port
      (u) - Untagged

VPN-ID   EVI-Name   EVI-Type Type Interface ESI              VLAN   DF-Status
Src-Addr Dst-Addr

100      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      3.3.3.3      ----
100      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      2.2.2.2      ----
100      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      4.4.4.4      ----
100      ----      --      AC   po1000.100 00:00:00:12:12:13:13:00:00:00 ----
      DF      ----      ----
101      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      3.3.3.3      ----
101      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      2.2.2.2      ----
101      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      4.4.4.4      ----
101      ----      --      AC   po1000.101 00:00:00:12:12:13:13:00:00:00 ----   NON-
      DF      ----      ----
102      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      3.3.3.3      ----
102      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      2.2.2.2      ----
102      ----      L2      NW   ----      ----              ----   ----
      1.1.1.1      4.4.4.4      ----
102      ----      --      AC   po1000.102 00:00:00:12:12:13:13:00:00:00 ----
      DF      ----      ----
Total number of entries are 12

```

Verify MPLS Tunnels Between PEs

- Ensures EVPN MPLS tunnels (per-EVI LSPs) between PEs are installed and operational.
- These tunnels provide underlay connectivity for E-Tree unicast and BUM forwarding.
- Tunnel status should be **Installed** for all peer PEs.
- Confirms label-switched connectivity between PEs, ensuring both Root and Leaf traffic can traverse the MPLS core.

```

PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination   Status      Up/Down      Update      evpn-id      Local-
Leaf Remote-Leaf Ext-Color FAT
=====
1.1.1.1      3.3.3.3      Installed   00:17:08     00:17:08     102          ---
      ---      ---
1.1.1.1      3.3.3.3      Installed   00:17:08     00:17:08     101          ---
      ---      ---
1.1.1.1      3.3.3.3      Installed   00:17:07     00:17:07     100          ---
      ---      ---
1.1.1.1      2.2.2.2      Installed   00:17:24     00:17:24     102          ---
      ---      ---
1.1.1.1      2.2.2.2      Installed   00:17:24     00:17:24     101          ---
      ---      ---
1.1.1.1      2.2.2.2      Installed   00:17:24     00:17:24     100          ---
      ---      ---
1.1.1.1      4.4.4.4      Installed   00:17:21     00:17:21     102          ---
      ---      ---
1.1.1.1      4.4.4.4      Installed   00:17:21     00:17:21     101          ---
      ---      ---

```

```

1.1.1.1      4.4.4.4      Installed      00:17:21      00:17:21      100      ---
---          ---          ---

Total number of entries are 9

```

```

PE1#show evpn mpls tunnel summary

Total number of entries: 9 [Installed: 9, Resolved: 0, Unresolved: 0]

```

Verify E-Tree Leaf Label Association

- Displays mapping of access interfaces to their E-Tree roles (Root or Leaf) and corresponding Leaf Label assignment.
- Root ACs show **Root** with no Leaf Label.
- Leaf ACs show **Leaf** with a unique Leaf Label (per-PE).
- Confirms correct Leaf identification per access interface and validates local Leaf Label programming used for unicast or BUM filtering.

```

PE1#show evpn etree-leaf brief

Leaf
Interface  Ifindex    Vnid      Leaf
           status  label
-----
po1000.100 0x1f400064 100       Root    ---
po1000.101 0x1f400065 101       Leaf    16
po1000.102 0x1f400066 102       Leaf    16

Total number of entries are 3

```

Verify Ethernet A-D per-ES (EAD) Route Advertisement

- Displays Ethernet Auto-Discovery (A-D) routes used for ESI discovery and Leaf Label distribution.
- In Scenario 2:
 - EAD routes with **ESI = 0** identify Leaf interfaces (single-homed or non-ESI).
 - EAD routes with **valid ESI values** represent multi-homed segments.
- Each Leaf AC advertises an A-D route with ESI = 0 and carries the assigned Leaf Label.
- Remote PEs receive these routes for BUM egress filtering.

RD[1.1.1.1:101]: ESI 0 VNID/LABEL 16 Nexthop 3.3.3.3

- Entries with ESI = 0 and Label = 16 correspond to Ethernet A-D per-ES (RT-1) routes carrying the Leaf Label.
- These routes confirm successful Leaf Label advertisement via BGP.
- The control-plane signaling for Leaf Label exchange is functioning, enabling:
 - Egress filtering for BUM traffic
 - Ingress filtering for unicast Leaf-to-Leaf communication

```

PE1#show bgp l2vpn evpn multihoming ethernet-ad-per-es
BGP table version is 7, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, a add-path, b back-up, * valid, > best, i -
internal,
              l - labeled, S Stale

```

Origin codes: i - IGP, e - EGP, ? - incomplete
 Description : Ext-Color - Extended community color

RD[1.1.1.1:100] VRF[vrf100]:

IP	ESI	Peer	Encap	Eth-Tag Flags	VNID/LABEL	Nexthop	
*							
i	00:00:00:12:12:13:13:00:00:00			4294967295	304	2.2.2.2	2.2.2.2 MPLS
*							
i	00:00:00:14:14:15:15:00:00:00			4294967295	304	4.4.4.4	4.4.4.4 MPLS
*							
i	00:00:00:14:14:15:15:00:00:00			4294967295	304	3.3.3.3	3.3.3.3 MPLS

RD[1.1.1.1:101] VRF[vrf101]:

IP	ESI	Peer	Encap	Eth-Tag Flags	VNID/LABEL	Nexthop	
*							
i	0			4294967295	16	3.3.3.3	3.3.3.3 MPLS
*							
i	0			4294967295	16	4.4.4.4	4.4.4.4 MPLS
*							
i	0			4294967295	16	2.2.2.2	2.2.2.2 MPLS
*							
i	00:00:00:12:12:13:13:00:00:00			4294967295	304	2.2.2.2	2.2.2.2 MPLS
*							
i	00:00:00:14:14:15:15:00:00:00			4294967295	304	4.4.4.4	4.4.4.4 MPLS
*							
i	00:00:00:14:14:15:15:00:00:00			4294967295	304	3.3.3.3	3.3.3.3 MPLS

RD[1.1.1.1:102] VRF[vrf102]:

IP	ESI	Peer	Encap	Eth-Tag Flags	VNID/LABEL	Nexthop	
*							
i	0			4294967295	16	3.3.3.3	3.3.3.3 MPLS
*							
i	0			4294967295	16	4.4.4.4	4.4.4.4 MPLS
*							
i	0			4294967295	16	2.2.2.2	2.2.2.2 MPLS
*							
i	00:00:00:12:12:13:13:00:00:00			4294967295	304	2.2.2.2	2.2.2.2 MPLS
*							
i	00:00:00:14:14:15:15:00:00:00			4294967295	304	4.4.4.4	4.4.4.4 MPLS
*							
i	00:00:00:14:14:15:15:00:00:00			4294967295	304	3.3.3.3	3.3.3.3 MPLS

RD[1.1.1.1:64512] VRF[evpn-gvrf-1]:

IP	ESI	Peer	Encap	Eth-Tag Flags	VNID/LABEL	Nexthop	
*>	0			4294967295	16	1.1.1.1	-----
		MPLS					
*>	00:00:00:12:12:13:13:00:00:00			4294967295	304	1.1.1.1	-----
		MPLS					

RD[2.2.2.2:64512]

IP	ESI	Peer	Encap	Eth-Tag Flags	VNID/LABEL	Nexthop	
*>i	0			4294967295	16	2.2.2.2	2.2.2.2 MP
LS							
*>i	00:00:00:12:12:13:13:00:00:00			4294967295	304	2.2.2.2	2.2.2.2 MP
LS							

RD[3.3.3.3:64512]

IP	ESI	Peer	Encap	Eth-Tag Flags	VNID/LABEL	Nexthop	
*>i	0			4294967295	16	3.3.3.3	3.3.3.3 MP
LS							
*>i	00:00:00:14:14:15:15:00:00:00			4294967295	304	3.3.3.3	3.3.3.3 MP
LS							

```
RD[4.4.4.4:64512]
      ESI              Eth-Tag      VNID/LABEL      Nexthop
IP      Peer          Encap        Flags
*>i    0              4294967295  16              4.4.4.4          4.4.4.4          MP
LS
*>i    00:00:00:14:14:15:15:00:00:00  4294967295  304              4.4.4.4          4.4.4.4          MP
LS
```

E-Tree Scenario 2 Network Topology Running Configurations

Here are the snippet configurations for all nodes in the given network topology [Figure 5](#).

Root PEs

PE1

```
PE1#show running-config

!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
!
hostname PE1-7038
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
feature rsyslog
logging remote server 10.16.100.20 5 port 1514 vrf management
lldp run
lldp tlv-select basic-mgmt port-description
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-capabilities
lldp tlv-select basic-mgmt system-description
lldp tlv-select basic-mgmt management-address
lldp notification-interval 1000
fault-management enable
!
evpn mpls enable
!
evpn esi hold-time 200
!
evpn etree enable scenario-2
!
evpn mpls multihoming enable
!
ip vrf management
!
mac vrf vrf100
rd 1.1.1.1:100
```

```
    route-target both 65010:100
!
mac vrf vrf101
  rd 1.1.1.1:101
  route-target both 65010:101
!
mac vrf vrf102
  rd 1.1.1.1:102
  route-target both 65010:102
!
evpn mpls vtep-ip-global 1.1.1.1
!
evpn mpls id 100
  host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
  host-reachability-protocol evpn-bgp vrf101
!
evpn mpls id 102
  host-reachability-protocol evpn-bgp vrf102
!
router ldp
  graceful-restart full
  targeted-peer ipv4 2.2.2.2
    exit-targeted-peer-mode
  targeted-peer ipv4 3.3.3.3
    exit-targeted-peer-mode
  targeted-peer ipv4 4.4.4.4
    exit-targeted-peer-mode
  transport-address ipv4 1.1.1.1
!
router rsvp
!
interface po10
  description connected-to-p1
  load-interval 30
  ip address 12.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface po20
  description connected-to-p2
  load-interval 30
  ip address 13.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface po1000
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.1212.1313
!
interface po1000.100 switchport
  encapsulation dot1q 100
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 100
!
interface po1000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
```

```
    map vpn-id 101 etree-leaf
!
interface po1000.102 switchport
 encapsulation dot1q 102
 rewrite push 0x8100 4090
 load-interval 30
 access-if-evpn
  map vpn-id 102 etree-leaf
!
interface ce5
 channel-group 1000 mode active
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/32 secondary
 ipv6 address ::1/128
 ip router isis ISIS-IGP
 enable-rsvp
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe10
 channel-group 10 mode active
!
interface xe11
 channel-group 10 mode active
!
interface xe12
 speed 10g
 channel-group 20 mode active
!
interface xe26
 channel-group 20 mode active
!
exit
!
router ospf 100
 ospf router-id 1.1.1.1
 bfd all-interfaces
 network 1.1.1.1/32 area 0.0.0.0
 network 12.1.1.0/24 area 0.0.0.0
 network 13.1.1.0/24 area 0.0.0.0
!
router bgp 65010
 neighbor EVPN peer-group
 neighbor EVPN remote-as 65010
 neighbor EVPN update-source lo
 neighbor EVPN advertisement-interval 0
 neighbor EVPN fall-over bfd multihop
 neighbor 2.2.2.2 peer-group EVPN
 neighbor 3.3.3.3 peer-group EVPN
 neighbor 4.4.4.4 peer-group EVPN
!
 address-family l2vpn evpn
  neighbor EVPN activate
 exit-address-family
!
exit
!
rsvp-path PE1-PE3 mpls
 13.1.1.2 strict
```

```

17.1.1.1 strict
19.1.1.2 strict
!
rsvp-path PE1-PE4 mpls
12.1.1.2 strict
17.1.1.2 strict
21.1.1.2 strict
!
rsvp-trunk PE1-PE3 ipv4
primary fast-reroute protection one-to-one
primary path PE1-PE3
to 3.3.3.3
!
rsvp-trunk PE1-PE4 ipv4
primary fast-reroute protection facility
primary path PE1-PE4
to 4.4.4.4
!
rsvp-trunk PE1-PE2 ipv4
to 2.2.2.2
!
line console 0
exec-timeout 0
!
!
end
!

```

PE2

```

PE2#show running-config

!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
hardware-profile port-config mode3
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE2-7039
port ce9 breakout 4X10g
port ce6 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt management-address
!
evpn mpls enable
!
evpn esi hold-time 200

```

```
!  
evpn etree enable scenario-2  
!  
evpn mpls multihoming enable  
!  
ip vrf management  
!  
mac vrf vrf100  
  rd 2.2.2.2:100  
  route-target both 65010:100  
!  
mac vrf vrf101  
  rd 2.2.2.2:101  
  route-target both 65010:101  
!  
mac vrf vrf102  
  rd 2.2.2.2:102  
  route-target both 65010:102  
!  
evpn mpls vtep-ip-global 2.2.2.2  
!  
evpn mpls id 100  
  host-reachability-protocol evpn-bgp vrf100  
!  
evpn mpls id 101  
  host-reachability-protocol evpn-bgp vrf101  
!  
evpn mpls id 102  
  host-reachability-protocol evpn-bgp vrf102  
!  
router ldp  
  graceful-restart full  
  targeted-peer ipv4 1.1.1.1  
    exit-targeted-peer-mode  
  targeted-peer ipv4 3.3.3.3  
    exit-targeted-peer-mode  
  targeted-peer ipv4 4.4.4.4  
    exit-targeted-peer-mode  
  transport-address ipv4 2.2.2.2  
!  
router rsvp  
!  
interface po30  
  description connected-to-pl  
  load-interval 30  
  ip address 14.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po40  
  load-interval 30  
  ip address 15.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po1000  
  switchport  
  load-interval 30  
  evpn multi-homed system-mac 0000.1212.1313  
!  
interface po1000.100 switchport  
  encapsulation dot1q 100  
  rewrite push 0x8100 4090  
  load-interval 30
```



```
access-if-evpn
  map vpn-id 100
!
interface po1000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
    map vpn-id 101 etree-leaf
!
interface po1000.102 switchport
  encapsulation dot1q 102
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
    map vpn-id 102 etree-leaf
!
interface ce6/1
  channel-group 40 mode active
!
interface ce6/2
  channel-group 40 mode active
!
interface ce6/3
  description connected-to-p2
  load-interval 30
  ip address 16.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce6/4
  description connected-to-p2
  load-interval 30
  ip address 16.1.2.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce7
!
interface ce8
  channel-group 1000 mode active
!
interface ce9/1
!
interface ce9/2
!
interface ce9/2.101 switchport
  encapsulation dot1q 101
  load-interval 30
  access-if-evpn
    map vpn-id 101 etree-leaf
!
interface ce9/2.102 switchport
  encapsulation dot1q 102
  load-interval 30
  access-if-evpn
    map vpn-id 102 etree-leaf
!
interface ce9/3
  channel-group 30 mode active
!
interface ce16
  channel-group 30 mode active
!
```

```
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 2.2.2.2/32 secondary
  ipv6 address ::1/128
  ip router isis ISIS-IGP
  enable-rsvp
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xe0
!
interface xe1
  speed 10g
!
exit
!
router ospf 100
  ospf router-id 2.2.2.2
  bfd all-interfaces
  network 2.2.2.2/32 area 0.0.0.0
  network 14.1.1.0/24 area 0.0.0.0
  network 15.1.1.0/24 area 0.0.0.0
  network 16.1.1.0/24 area 0.0.0.0
  network 16.1.2.0/24 area 0.0.0.0
!
router bgp 65010
  neighbor EVPN peer-group
  neighbor EVPN remote-as 65010
  neighbor EVPN update-source lo
  neighbor EVPN advertisement-interval 0
  neighbor EVPN fall-over bfd multihop
  neighbor 1.1.1.1 peer-group EVPN
  neighbor 3.3.3.3 peer-group EVPN
  neighbor 4.4.4.4 peer-group EVPN
!
  address-family l2vpn evpn
  neighbor EVPN activate
  exit-address-family
!
exit
!
rsvp-path PE2-PE3 mpls
  15.1.1.2 loose
  17.1.1.1 loose
  19.1.1.2 loose
!
rsvp-path PE2-PE4 mpls
  14.1.1.2 loose
  17.1.1.2 loose
  21.1.1.2 loose
!
rsvp-trunk PE2-PE3 ipv4
  primary fast-reroute protection one-to-one
  primary path PE2-PE3
  to 3.3.3.3
!
rsvp-trunk PE2-PE4 ipv4
  primary fast-reroute protection facility
  primary path PE2-PE4
  to 4.4.4.4
!
```

```

rsvp-trunk PE2-PE1 ipv4
to 1.1.1.1
!
line console 0
exec-timeout 0
!
!
end

!

```

Leaf PEs

PE3

```

PE3#show running-config
!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE3-7032
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature ssh vrf management
feature dns relay
ip dns relay
ipv6 dns relay
feature ntp vrf management
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt management-address
!
evpn mpls enable
!
evpn etree enable scenario-2
!
evpn mpls multihoming enable
!
ip vrf management
!
mac vrf vrf100
rd 3.3.3.3:100
route-target both 65010:100
!
mac vrf vrf101
rd 3.3.3.3:101
route-target both 65010:101
!
mac vrf vrf102
rd 3.3.3.3:102
route-target both 65010:102
evpn mpls vtep-ip-global 3.3.3.3
!

```

```
evpn mpls id 100
  host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
  host-reachability-protocol evpn-bgp vrf101
!
evpn mpls id 102
  host-reachability-protocol evpn-bgp vrf102
!
router ldp
  graceful-restart full
  targeted-peer ipv4 1.1.1.1
    exit-targeted-peer-mode
  targeted-peer ipv4 2.2.2.2
    exit-targeted-peer-mode
  targeted-peer ipv4 4.4.4.4
    exit-targeted-peer-mode
  transport-address ipv4 3.3.3.3
!
router rsvp
!
interface po2000
  switchport
  load-interval 30
  evpn multi-homed system-mac 0000.1414.1515
!
interface po2000.100 switchport
  encapsulation dot1q 100
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 100
!
interface po2000.101 switchport
  encapsulation dot1q 101
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 101 etree-leaf
!
interface po2000.102 switchport
  encapsulation dot1q 102
  rewrite push 0x8100 4090
  load-interval 30
  access-if-evpn
  map vpn-id 102 etree-leaf
!
interface sa10
  description connected-to-p1
  load-interval 30
  ip address 19.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface sa20
  description connected-to-p2
  load-interval 30
  ip address 20.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
```

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 3.3.3.3/32 secondary  
  ipv6 address ::1/128  
  ip router isis ISIS-IGP  
  enable-rsvp  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe4  
!  
interface xe5  
  speed lg  
  channel-group 2000 mode active  
!  
interface xe8.101 switchport  
  encapsulation dot1q 101  
  load-interval 30  
  access-if-evpn  
  map vpn-id 101  
!  
interface xe22  
  static-channel-group 20  
!  
interface xe23  
  static-channel-group 20  
!  
interface xe24  
  static-channel-group 10  
!  
interface xe25  
  static-channel-group 10  
!  
exit  
!  
router ospf 100  
  ospf router-id 3.3.3.3  
  bfd all-interfaces  
  network 3.3.3.3/32 area 0.0.0.0  
  network 19.1.1.0/24 area 0.0.0.0  
  network 20.1.1.0/24 area 0.0.0.0  
!  
router bgp 65010  
  neighbor EVPN peer-group  
  neighbor EVPN remote-as 65010  
  neighbor EVPN update-source lo  
  neighbor EVPN advertisement-interval 0  
  neighbor EVPN fall-over bfd multihop  
  neighbor 1.1.1.1 peer-group EVPN  
  neighbor 2.2.2.2 peer-group EVPN  
  neighbor 4.4.4.4 peer-group EVPN  
!  
  address-family l2vpn evpn  
  neighbor EVPN activate  
  exit-address-family  
!  
exit  
!  
  rsvp-path PE3-PE1 mpls  
    20.1.1.1 strict  
    17.1.1.1 strict  
    12.1.1.1 strict  
  !  
  rsvp-path PE3-PE2 mpls
```

```
19.1.1.1 strict
17.1.1.2 strict
15.1.1.1 strict
!
rsvp-trunk PE3-PE1 ipv4
primary fast-reroute protection facility
primary path PE3-PE1
to 1.1.1.1
!
rsvp-trunk PE3-PE2 ipv4
primary fast-reroute protection one-to-one
primary path PE3-PE2
to 2.2.2.2
!
rsvp-trunk PE3-PE4 ipv4
to 4.4.4.4
!
line console 0
exec-timeout 0
!
!
end
!
```

PE4

```
PE4#show running-config

!
hardware-profile filter ingress-ipv6-ext-subif enable
hardware-profile filter ingress-ipv4-subif enable
hardware-profile filter egress-ipv4-ext enable
hardware-profile filter egress-ipv6 enable
hardware-profile filter evpn-mpls-mh enable
hardware-profile filter twamp-ipv4 enable
hardware-profile statistics voq-full-color enable
hardware-profile statistics cfm-ccm enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname PE4-7033
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
evpn mpls arp-nd refresh-timer 120
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt management-address
!
evpn mpls enable
!
evpn esi hold-time 200
!
evpn etree enable scenario-2
!
evpn mpls multihoming enable
!
ip vrf management
!
mac vrf vrf100
```

```
rd 4.4.4.4:100
route-target both 65010:100
!
mac vrf vrf101
rd 4.4.4.4:101
route-target both 65010:101
!
mac vrf vrf102
rd 4.4.4.4:102
route-target both 65010:102
!
evpn mpls vtep-ip-global 4.4.4.4
!
evpn mpls mac-ageing-time 360
!
evpn mpls id 100
host-reachability-protocol evpn-bgp vrf100
!
evpn mpls id 101
host-reachability-protocol evpn-bgp vrf101
!
evpn mpls id 102
host-reachability-protocol evpn-bgp vrf102
!
router ldp
graceful-restart full
targeted-peer ipv4 1.1.1.1
exit-targeted-peer-mode
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
transport-address ipv4 4.4.4.4
!
router rsvp
!
interface po2000
switchport
load-interval 30
evpn multi-homed system-mac 0000.1414.1515
!
interface po2000.100 switchport
encapsulation dot1q 100
rewrite push 0x8100 4090
load-interval 30
access-if-evpn
map vpn-id 100
!
interface po2000.101 switchport
encapsulation dot1q 101
rewrite push 0x8100 4090
load-interval 30
access-if-evpn
map vpn-id 101 etree-leaf
!
interface po2000.102 switchport
encapsulation dot1q 102
rewrite push 0x8100 4090
load-interval 30
access-if-evpn
map vpn-id 102 etree-leaf
!
interface sa30
description connected-to-p1
load-interval 30
ip address 18.1.1.2/24
mtu 9216
label-switching
```

```
enable-ldp ipv4
enable-rsvp
!
interface sa40
description connected-to-p2
load-interval 30
ip address 21.1.1.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface cd0
!
interface cd1
!
interface ce2
!
interface ce3
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
interface lo
ip address 127.0.0.1/8
ip address 4.4.4.4/32 secondary
ipv6 address ::1/128
ip router isis ISIS-IGP
enable-rsvp
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xe6
speed 10g
!
interface xe7
!
interface xe8
speed 10g
static-channel-group 30
!
interface xe9
speed 1g
channel-group 2000 mode active
!
interface xe10
!
interface xe11
static-channel-group 40
!
interface xe24
static-channel-group 30
!
exit
!
router ospf 100
ospf router-id 4.4.4.4
bfd all-interfaces
network 4.4.4.4/32 area 0.0.0.0
network 18.1.1.0/24 area 0.0.0.0
network 21.1.1.0/24 area 0.0.0.0
!
router bgp 65010
neighbor EVPN peer-group
```



```

neighbor EVPN remote-as 65010
neighbor EVPN update-source lo
neighbor EVPN advertisement-interval 0
neighbor EVPN fall-over bfd multihop
neighbor 1.1.1.1 peer-group EVPN
neighbor 2.2.2.2 peer-group EVPN
neighbor 3.3.3.3 peer-group EVPN
!
address-family l2vpn evpn
neighbor EVPN activate
exit-address-family
!
exit
!
rsvp-path PE4-PE1 mpls
21.1.1.1 strict
13.1.1.1 strict
!
rsvp-path PE4-PE2 mpls
18.1.1.1 strict
14.1.1.1 strict
!
rsvp-trunk PE4-PE1 ipv4
primary fast-reroute protection one-to-one
primary path PE4-PE1
to 1.1.1.1
!
rsvp-trunk PE4-PE2 ipv4
primary fast-reroute protection facility
primary path PE4-PE2
to 2.2.2.2
!
rsvp-trunk PE4-PE3 ipv4
to 3.3.3.3
!
line console 0
exec-timeout 0
!
!
end
!

```

Core Routers

P1

```

P1#show running-config
!
hardware-profile statistics ingress-acl enable
!
bfd interval 3 minrx 3 multiplier 3
!
qos enable
!
hostname P1-7034
port ce49 breakout 4X10g
ip name-server vrf management 10.16.10.23
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
lldp run
lldp tlv-select basic-mgmt system-name

```

```
lldp tlv-select basic-mgmt management-address
!
ip vrf management
!
router ldp
  transport-address ipv4 5.5.5.5
!
router rsvp
!
interface po10
  load-interval 30
  ip address 12.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface po30
  description connected-to-pe2
  load-interval 30
  ip address 14.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface sa10
  description connected-to-pe3
  load-interval 30
  ip address 19.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface sa30
  description connected-to-pe4
  load-interval 30
  ip address 18.1.1.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce49/1
  static-channel-group 10
!
interface ce49/2
  static-channel-group 10
!
interface ce49/3
  static-channel-group 30
!
interface ce49/4
  description connected-to-p2
  load-interval 30
  ip address 17.1.2.1/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce52
  channel-group 30 mode active
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
```

```
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 5.5.5.5/32 secondary  
  ipv6 address ::1/128  
  ip router isis ISIS-IGP  
  enable-rsvp  
!  
interface lo.management  
  ip vrf forwarding management  
  ip address 127.0.0.1/8  
  ipv6 address ::1/128  
!  
interface xe1  
  description connected-to-p2  
  load-interval 30  
  ip address 17.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface xe2  
  channel-group 10 mode active  
!  
interface xe3  
  channel-group 10 mode active  
!  
interface xe4  
!  
interface xe5  
  static-channel-group 30  
!  
interface xe45  
  channel-group 30 mode active  
!  
  exit  
!  
router ospf 100  
  ospf router-id 5.5.5.5  
  bfd all-interfaces  
  network 5.5.5.5/32 area 0.0.0.0  
  network 12.1.1.0/24 area 0.0.0.0  
  network 14.1.1.0/24 area 0.0.0.0  
  network 17.1.1.0/24 area 0.0.0.0  
  network 17.1.2.0/24 area 0.0.0.0  
  network 18.1.1.0/24 area 0.0.0.0  
  network 19.1.1.0/24 area 0.0.0.0  
!  
line console 0  
  exec-timeout 0  
!  
!  
end  
!
```

P2

```
P2#show running-config  
!  
hardware-profile statistics ingress-acl enable  
!  
bfd interval 3 minrx 3 multiplier 3  
!  
qos enable
```

```
!  
hostname P2-7035  
port ce49 breakout 4X10g  
port ce50 breakout 4X10g  
ip name-server vrf management 10.16.10.23  
tfo Disable  
errdisable cause stp-bpdu-guard  
feature dns relay  
ip dns relay  
ipv6 dns relay  
feature rsyslog  
logging remote server 10.16.100.20 5 port 1514 vrf management  
lldp run  
lldp tlv-select basic-mgmt port-description  
lldp tlv-select basic-mgmt system-name  
lldp tlv-select basic-mgmt system-capabilities  
lldp tlv-select basic-mgmt system-description  
lldp tlv-select basic-mgmt management-address  
lldp notification-interval 1000  
fault-management enable  
!  
ip vrf management  
!  
router ldp  
  transport-address ipv4 6.6.6.6  
!  
router rsvp  
!  
interface po20  
  description connected-to-pe1  
  load-interval 30  
  ip address 13.1.1.2/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface po40  
  description connected-to-pe2  
  load-interval 30  
  ip address 15.1.1.2/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface sa20  
  description connected-to-pe3  
  load-interval 30  
  ip address 20.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface sa40  
  description connected-to-pe4  
  load-interval 30  
  ip address 21.1.1.1/24  
  mtu 9216  
  label-switching  
  enable-ldp ipv4  
  enable-rsvp  
!  
interface ce49/1  
  channel-group 40 mode active  
!  
interface ce49/2
```

```
    channel-group 40 mode active
!
interface ce49/3
  description connected-to-pe2
  load-interval 30
  ip address 16.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce49/4
  description connected-to-pe2
  load-interval 30
  ip address 16.1.2.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface ce50/1
  static-channel-group 20
!
interface ce50/2
!
interface ce50/3
  static-channel-group 20
!
interface ce50/4
  channel-group 20 mode active
!
interface eth0
  ip vrf forwarding management
  ip address dhcp
!
interface lo
  ip address 127.0.0.1/8
  ip address 6.6.6.6/32 secondary
  ipv6 address ::1/128
  ip router isis ISIS-IGP
  enable-rsvp
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface xe6
  description connected-to-pl
  load-interval 30
  ip address 17.1.1.2/24
  mtu 9216
  label-switching
  enable-ldp ipv4
  enable-rsvp
!
interface xe7
  channel-group 20 mode active
!

interface xell
  static-channel-group 40
!
interface xe20
  label-switching
!
interface xe47
  description connected-to-pl
```

```
load-interval 30
ip address 17.1.2.2/24
mtu 9216
label-switching
enable-ldp ipv4
enable-rsvp
!
interface xe48
!
exit
!
router ospf 100
ospf router-id 6.6.6.6
bfd all-interfaces
network 6.6.6.6/32 area 0.0.0.0
network 13.1.1.0/24 area 0.0.0.0
network 15.1.1.0/24 area 0.0.0.0
network 16.1.1.0/24 area 0.0.0.0
network 16.1.2.0/24 area 0.0.0.0
network 17.1.1.0/24 area 0.0.0.0
network 17.1.2.0/24 area 0.0.0.0
network 20.1.1.0/24 area 0.0.0.0
network 21.1.1.0/24 area 0.0.0.0
!
line console 0
exec-timeout 0
!
!
end
!
```

Access Switches

SW-1

```
SW1#show running-config
!
ip vrf management
!
hostname SW1-7037
ip domain-lookup
bridge 1 protocol ieee vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
feature telnet vrf
no feature telnet
feature ssh vrf
no feature ssh
snmp-server enable snmp vrf
snmp-server view all .1 included vrf
feature ntp vrf
ntp enable vrf
feature rsyslog vrf
lldp run
lldp tlv-select basic-mgmt system-name
lldp tlv-select basic-mgmt system-description
!
vlan database
vlan-reservation 4000-4062
vlan 2-1000 bridge 1 state enable
!
interface po1000
switchport
```

```
bridge-group 1 spanning-tree disable
switchport mode trunk
switchport trunk allowed vlan all
!
interface lo
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface lo.management
ip vrf forwarding management
ip address 127.0.0.1/8
ipv6 address ::1/128
!
interface xel/1
!
interface xel/2
!
interface xel/3
!
interface xel/4
!
interface xe2
!
interface xe3/1
port breakout enable 4X10g
switchport
bridge-group 1 spanning-tree disable
switchport mode trunk
switchport trunk allowed vlan all
!
interface xe3/2
!
interface xe5/1
channel-group 1000 mode active
!
interface xe8
channel-group 1000 mode active
!
no mac-address-table learning bridge 1 interface po1000
no mac-address-table learning bridge 1 interface xe3/1
!
line console 0
exec-timeout 0
!
!
end
```

E-Tree Scenario 2 Implementation Examples

In a real-world enterprise WAN or broadband service scenario, multiple branch offices connect to regional PE devices. Using OcNOS EVPN E-Tree Scenario 2, the provider ensures that branches (Leaf sites) cannot communicate directly with each other, enforcing strict isolation. Data centers (Spine sites) continue to communicate with all branches and other Root sites. EVPN E-Tree Scenario 2 uses Leaf Labels, ingress unicast filtering, and egress BUM filtering to enforce these rules efficiently, even on PE devices hosting both Root and Leaf sites. This design supports both single-homed and multi-homed branches, providing a secure, scalable, and high-performance solution for multi-tenant or enterprise deployments.

E-Tree Scenario 2 Troubleshooting

- Missing RT-2 route attributes: Verify per-AC or EVI configuration.

- Leaf Label not programmed: Check route reception and hardware resources.
 - Unexpected Leaf-to-Leaf forwarding: Confirm Leaf roles and ingress filtering rules.
 - BUM flooding to Leaf ACs: Check Leaf Label advertisement in Ethernet A-D routes.
 - CLI not active: Ensure `evpn etree enable scenario-2` command is configured.
-

E-Tree Scenario 2 Revised Commands

- Introduced the scenario-2 parameter in the [evpn etree \(page 209\)](#) command to enable EVPN MPLS E-Tree SC-2, which configures the device to support both Root and Leaf sites on the same PE device.
- Added `etree-leaf` parameter in the [map vpn-id \(page 210\)](#) to enable configuring the access interface (AC) or sub-interface as a Leaf in an EVPN E-Tree deployment.

evpn etree

Use this command to enable E-Tree functionality within the EVPN configuration. The `evpn etree enable` option ([EVPN MPLS E-Tree Scenario 1](#)) configures the device to function as either a Leaf or a Root site per EVPN Instance (EVI), as defined in RFC-8317. The `scenario-2` option ([EVPN MPLS E-Tree Scenario 2 \(page 177\)](#)) configures the device to support both Root and Leaf sites on the same Provider Edge (PE) device, enforcing traffic filtering rules to prevent Leaf-to-Leaf communication while allowing Root-to-Root and Root-to-Leaf traffic.



Note: If Scenario 1 is configured, it must be removed before enabling Scenario 2.

Command Syntax

```
evpn etree enable (scenario-2|)
no evpn etree enable (scenario-2|)
```

Parameters

scenario-2

Enables EVPN MPLS E-Tree Scenario 2, which configures the device to support both Root and Leaf sites on the same PE device.

Default

Disabled

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 6.5.1. Introduced the `scenario-2` parameter in OcNOS version 7.0.0.

Example

The following example illustrates how to activate E-Tree Scenario 1 functionality for EVPN, which configures the device as either a Leaf or a Root site per EVI.

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable

OcNOS(config)#no evpn etree enable
```

The following example illustrates how to activate E-Tree Scenario 2 functionality for EVPN.

```
OcNOS#configure terminal
OcNOS(config)#evpn etree enable scenario-2
```

map vpn-id

Use this command to map a sub-interface to a tenant.

Use the no form of this command to remove the tenant.

Command Syntax

```
map vpn-id <1-16777215> (etree-leaf|)  
no map vpn-id <1-16777215>
```

Parameters

vpn-id <1-16777215>

Specifies the EVPN Instance (EVI) ID (VNID) to which the interface belongs.

etree-leaf

(Optional) Enables configuring the access interface (AC) or sub-interface as a Leaf in an EVPN E-Tree deployment. This allows the PE to enforce E-Tree traffic rules, including Leaf-to-Leaf traffic blocking and Root-to-Leaf communication.

Command Mode

Access interface EVPN mode

Applicability

Introduced in OcNOS version 3.0. Added `etree-leaf` parameter in OcNOS version 7.0.0.

Examples

1. Configure a regular EVPN interface (no E-Tree Leaf):

```
#configure terminal  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#map vpn-id 1  
(config-access-if)#end
```

2. Configure an interface as an E-Tree Leaf AC:

```
#configure terminal  
(config)#evpn etree enable scenario-2  
(config)#evpn mpls id 203 etree-leaf  
(config-evpn-mpls)#host-reachability-protocol evpn-bgp vrf103  
(config-evpn-mpls)#exit  
  
(config)#interface xel.1 switchport  
(config-if)#access-if-evpn  
(config-access-if)#map vpn-id 3 etree-leaf  
(config-access-if)#end
```

E-Tree Scenario 2 Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms or Acronym	Description
Attachment Circuit (AC)	Physical or logical interface (e.g., sub-interface, port-channel) connecting a customer VLAN or segment to a PE router.
Broadcast, Unknown-unicast, Multicast (BUM)	Collective term for Layer 2 traffic types that require flooding or replication across EVPN sites.
Designated Forwarder (DF)	PE router elected to forward BUM traffic on a multi-homed Ethernet Segment (ESI). Prevents duplication.
Ethernet Auto-Discovery (EAD)	EVPN route type used for signaling Ethernet Segment (ES) membership and local Leaf Label advertisement.
E-Tree	Ethernet VPN service model that defines hierarchical connectivity between Root and Leaf sites, restricting Leaf-to-Leaf communication.
Ethernet Segment (ES)	Group of physical links connecting a customer device (CE or switch) redundantly to two or more PEs, identified by an ESI.
Ethernet Segment Identifier (ESI)	10-byte unique identifier representing a multi-homed access link in EVPN.
Ethernet VPN (EVPN)	BGP-based VPN technology providing Layer 2 services over MPLS or IP networks.
EVPN Instance (EVI)	Logical service instance that defines the scope of MAC learning and forwarding for a given EVPN.
Interior Gateway Protocol (IGP)	Routing protocol (e.g., IS-IS or OSPF) used to exchange underlay routes among core and PE routers.
Label-Switching	MPLS data-plane forwarding method that uses labels to route packets efficiently through the network.
Leaf Label	Locally assigned label identifying a Leaf access circuit; used to enforce Leaf-to-Root communication only.
Label Distribution Protocol (LDP)	MPLS control protocol that distributes label bindings for IGP prefixes.
Multiprotocol Label Switching (MPLS)	Transport technology that uses short labels instead of IP lookups to forward packets through the network.
Multi-Homed Provider Edge (MH-PE)	PE router that connects to a customer site via redundant Ethernet Segments shared with another PE.
Provider Edge (PE)	Edge router that connects customer networks (CEs or switches) to the provider MPLS backbone.
Provider (P) Router	Core router within the MPLS network that forwards labeled traffic but does not maintain customer VPN routes.
Resource Reservation Protocol (RSVP)	MPLS signaling protocol that establishes label-switched paths (LSPs) with traffic engineering (TE) attributes.
Root AC	Access circuit configured as Root in an E-Tree service; can send to or receive from any Leaf.
Leaf AC	Access circuit configured as Leaf in an E-Tree service; restricted from sending traffic directly to another Leaf.
Scenario-2 (SC-2)	Enhanced E-Tree mode in OcNOS where Leaf Label is advertised via BGP to achieve ingress or egress filtering.

Key Terms or Acronym	Description
Virtual Local Area Network (VLAN)	Logical network segment used to separate broadcast domains within Layer 2 infrastructure.
VPN-ID	Identifier that maps a set of interfaces or EVPN instances to a common Layer 2 or Layer 3 VPN service.
Route Target-2 (RT-2)	MAC or IP advertisement route carrying E-Tree extended community and Leaf Label.
Route Target-1 (RT-1)	Ethernet A-D per ES route carrying Leaf Label information.
Split Horizon	Prevents Leaf-to-Leaf forwarding within the same PE.

NETWORK MANAGEMENT AND AUTOMATION

OcNOS advances its programmability and management capabilities with extended NetConf or YANG models and enhanced visibility. These features enable easier automation, monitoring, and integration with modern network orchestration systems.

Container Runtime and Life cycle Management Using K3S	214
Overview	214
Prerequisites	215
Configuration	215
Implementation Examples	217
K3s Containers Commands	220
Glossary	222
NetConf Access Control Model User Guide	224
Overview	224
Feature Characteristics	224
Configuration	225
Common NACM Rule Fields	226
Creating NetConf RPC for NACM	229
Glossary	238
sFlow - Sample Packet Monitoring for Multiple Interfaces	239
Overview	239
Features Characteristics	240
Benefits	240
Prerequisites	240
Configuration	241
Configuring sFlow with User Defined VRFs	243
Implementation Examples	246
Commands	252
Troubleshooting	255
VxLAN OAM for Overlay Networks	256
CLI-Script and CLI-Shell	257
System Limits and Counters	274
Overview	274
System Limits and Counters Limitation	275
System Limits and Counters Configuration	275
System Limits and Counters Implementation Example	282
System Limits and Counters Commands	282
System Limit Counters Troubleshooting	297
System Limit Counters Glossary	298

Container Runtime and Life cycle Management Using K3S

Overview

K3S is a lightweight Kubernetes distribution integrated into OcNOS, turning switch/router into a Edge Compute Platform. It allows the network device to natively host and manage containerized third-party applications, such as NetConf servers or telemetry collectors and running Kubernetes Pods directly inside the Network OS environment. K3s uses its built-in runtime, containerd, to handle all low-level container tasks like pulling images. The K3s system actively manages the container's lifecycle, ensuring the application always runs. This orchestration creates the containers, monitor the health, and scaling or healing failed containers and keeps the service running smoothly on the node.

Feature Characteristics

- **Core Integration & Purpose**
 - Built-in Orchestration: Integrates a lightweight Kubernetes (K3s) engine directly into the Network OS.
 - Third-Party Hosting: Enables the system (switch/router) to host, manage, and run third-party containerized applications (Pods).
 - Eliminates External Servers: Allows services like NetConf servers, SNMP daemons, or telemetry collectors to run as local Pods, removing the need for dedicated external compute servers.
- **Networking & Connectivity**
 - Pod Networking: Utilizes Linux networking primitives: Linux namespaces (`zebosfib0` & `zebosfib1`), veth pairs (`fib0veth-fib1veth`), and a CNI bridge (`cni0`) for internal Pod communication.
 - External Access (NAT): External clients access Pods via the OcNOS management IP using NAT/port-forwarding.
 - Internal Communication (Direct): OcNOS internal processes communicate with Pods directly using IPv4 routes.
- **Management & Control**
 - Resource Control: Provides fine-grained resource control (CPU, memory, storage) for hosted Pods.
 - Observability: Supports standard K3s metrics for monitoring and health checks.
 - Security: Features auto-managed certificates for securing communications.
 - Operational Model: Effectively turns OcNOS into a network OS with built-in container orchestration.

Benefits

The K3s feature lets users run third-party or custom applications directly on the network device, eliminating the need for external servers.

- **Integration:** Deploy NetConf, SNMP, telemetry, or RCA tools as local Pods on OcNOS.
- **Efficiency:** Saves hardware, power, and maintenance by consolidating services on one platform.
- **Automation:** Uses Kubernetes orchestration for easy Pod life cycle management.
- **Flexibility:** Supports any containerized app with defined CPU, memory, and storage limits.
- **Visibility:** Built-in metrics for resource usage and health monitoring.

- **Scalability:** Simplifies adding or upgrading services without OS modification.

Prerequisites

The following conditions must be satisfied before enabling and using K3s on OcNOS:

- Ensure OcNOS has K3s-enabled image, configured `namespaces/veth` pair, Network Address Translation (NAT) rules, and required YAML/image files before enabling the feature.
- **System Requirements**
 - OcNOS image must include K3s support (from version supporting container integration).
 - Should be sufficient CPU (≥ 2 cores) and Memory (≥ 2 GB) available for K3s and Pods.
 - Required adequate storage in `/cfg` or `/var/lib/rancher/k3s` for images and manifests.
- **Network Setup**
 - Properly configured management (`eth0`) interface in `zebosfib1` namespace.

Limitations

- **Single-node only:** K3s runs locally on OcNOS; no multi-node cluster support.
- **Resource-bound:** Limited by device CPU, memory, and storage — heavy containers may impact NOS performance.
- **No Internet registry access:** Images must be loaded manually (offline).
- **Limited Pod count:** Only a few Pods can run reliably due to hardware constraints.
- **Static NAT/ports:** External access requires predefined or manually added port-forwarding rules.
- **No GUI dashboard:** Management via CLI (`kubectl`, `k3s ctr`) only.
- **Basic persistence:** Storage limited to host paths (e.g., `/cfg`) — no dynamic storage provisioning.

Configuration

Enable K3s

```
OcNOS#exec-shell systemctl start k3s.service
```

Container images — import into k3s

```
OcNOS#exec-shell k3s ctr image import netconfd-v1.tar
```

Pod & Service manifests install

```
OcNOS#exec-shell kubectl apply -f netconfd-pod.yaml
OcNOS#exec-shell kubectl apply -f netconfd-svc-lb.yaml
```

Verify & monitor

```
OcNOS#exec-shell kubectl get pods -o wide
OcNOS#exec-shell kubectl get svc -o wide
OcNOS#exec-shell kubectl get nodes -o wide
```

Topology

Single Node

Validation

OcNOS#exec-shell kubectl describe netconfd-pod

```
OcNOS#exec-shell kubectl describe pod netconfd-pod
Name:          netconfd-pod
Namespace:     default
Priority:       0
Service Account: default
Node:          ocnos-node/240.0.0.2
Start Time:    Thu, 23 Oct 2025 16:53:23 +0000
Labels:        app=netconfd
Annotations:   <none>
Status:        Running
IP:            203.1.114.6
IPs:
  IP: 203.1.114.6
Containers:
  netconfd-cnt:
    Container
ID:   containerd://752462ea4ed329ac4670ef63321d86ece3e713d09532d7108ed56926257440b4
Image:          netconfd:v1
Image ID:       sha256:1e6d89f15f8d1cblca97733c7053bcd1df5aa9959012e7f985ee0dd9d6690e90
Port:           <none>
Host Port:      <none>
State:          Running
  Started:      Thu, 23 Oct 2025 16:53:24 +0000
Ready:          True
Restart Count:  0
Environment:    <none>
Mounts:
  /var/run/secrets/kubernetes.io/serviceaccount from kube-api-access-vpp28 (ro)
Conditions:
  Type                               Status
PodReadyToStartContainers            True
Initialized                          True
Ready                                True
ContainersReady                      True
PodScheduled                         True
Volumes:
  kube-api-access-vpp28:
    Type:          Projected (a volume that contains injected data from multiple
sources)
    TokenExpirationSeconds: 3607
    ConfigMapName:      kube-root-ca.crt
    Optional:           false
    DownwardAPI:        true
QoS Class:           BestEffort
Node-Selectors:      <none>
Tolerations:         node.kubernetes.io/not-ready:NoExecute op=Exists for 300s
                     node.kubernetes.io/unreachable:NoExecute op=Exists for 300s
Events:
  Type    Reason      Age   From          Message
  ----    -
Normal   Scheduled   27m   default-scheduler   Successfully assigned default/netconfd-pod to
ocnos-node
Normal   Pulled      27m   kubelet         Container image "netconfd:v1" already present on
machine
Normal   Created     27m   kubelet         Created container: netconfd-cnt
Normal   Started     27m   kubelet         Started container netconfd-cnt
OcNOS#
```


OcNOS#exec-shell kubectl describe svc netconfd-svc-lb

```

Name:                netconfd-svc-lb
Namespace:           default
Labels:              app=netconfd
Annotations:         <none>
Selector:            app=netconfd
Type:                LoadBalancer
IP Family Policy:    SingleStack
IP Families:         IPv4
IP:                  203.0.114.215
IPs:                 203.0.114.215
LoadBalancer Ingress: 240.0.0.2 (VIP)
Port:                <unset> 10830/TCP
TargetPort:          830/TCP
NodePort:            <unset> 32275/TCP
Endpoints:           203.1.114.6:830
Session Affinity:    None
External Traffic Policy: Cluster
Internal Traffic Policy: Cluster
Events:
  Type    Reason              Age   From                      Message
  ----    -
  Normal  EnsuringLoadBalancer  28m   service-controller        Ensuring load balancer
  Normal  AppliedDaemonSet     28m   service-lb-controller     Applied LoadBalancer DaemonSet
kube-system/svclb-netconfd-svc-lb-17b17ff4
  Normal  UpdatedLoadBalancer  28m   service-lb-controller     Updated LoadBalancer with new IPs:
[] -> [240.0.0.2]
OcNOS#

```

Implementation Examples

```

netconfd-pod.yaml
---
apiVersion: v1
kind: Pod
metadata:
  name: netconfd-pod
  labels:
    app: netconfd
spec:
  containers:
  - name: netconfd-cnt
    image: netconfd:v1
    imagePullPolicy: Never

```

```

netconfd-svc-lb.yaml
---
apiVersion: v1
kind: Service
metadata:
  name: netconfd-svc-lb
  labels:
    app: netconfd
spec:
  selector:
    app: netconfd
  ports:
  - protocol: TCP
    port: 10830
    targetPort: 830
  type: LoadBalancer

```

```

OcNOS#pwd
/root

```

```

OcNOS#cd /home/ocnos
OcNOS#ls -ltr
total 118816
-rw-r----- 1 ocnos root 121658880 Apr 25 13:33 netconfd-v1.tar
-rw-r----- 1 ocnos root      215 Apr 25 13:33 netconfd-svc-lb.yaml
-rw-r----- 1 ocnos root      181 Apr 25 13:33 netconfd-pod.yaml

OcNOS#exec-shell k3s ctr image import netconfd-v1.tar
docker.io/library/netconfd:v1 saved
application/vnd.oci.image.manifest.v1+json
sha256:799d5ce61ed32420a80afd9fb027e8932b210b82dffc4cbaa6acbaefe5c74c77
Importing elapsed: 8.0 s total: 0.0 B (0.0 B/s)

OcNOS#exec-shell k3s ctr image ls
REF                                     TYPE                                SIZE    PLATF
ORMS  LABELS
docker.io/library/netconfd:v1          application/vnd.oci.image.manifest.v1+json sha256:799d5ce61ed32420a80afd9fb027e8932b210b82dffc4cbaa6acbaefe5c74c77 116.0 MiB linux/amd64 io.cri-containerd.image=managed
docker.io/rancher/klipper-helm:v0.9.8-build20250709 application/vnd.oci.image.manifest.v1+json sha256:ef2a7527aea909ee1299efccbad971b7f86e8682311d07970972f57ac6f3e9b5 206.0 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/klipper-lb:v0.4.13 application/vnd.oci.image.manifest.v1+json sha256:89a128a64d0f07434abddf496ce78d20208accc263414668bb6d7a978c220226 12.1 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/local-path-provisioner:v0.0.31 application/vnd.oci.image.manifest.v1+json sha256:74de96b9971dd4ea1a835687b80da0619820b6eb41f6f8f8032bdcfe05e791db 57.9 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/mirrored-coredns-coredns:1.12.3 application/vnd.oci.image.manifest.v1+json sha256:7a6fa2b0e04fb718f81c554cd569ae76dca021205551c20e6falf2f0a2106870 71.7 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/mirrored-library-busybox:1.36.1 application/vnd.oci.image.manifest.v1+json sha256:21a5bc93624c80bf8022b7bfad82a3242fa380bb9869300fe5ab4326311ae0ff 4.3 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/mirrored-library-traefik:3.3.6 application/vnd.oci.image.manifest.v1+json sha256:769d5df52751c361a149921905c7dff87f5c3033e8118883e7368d60ffaabc91 214.5 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/mirrored-metrics-server:v0.8.0 application/vnd.oci.image.manifest.v1+json sha256:4557c70a82b0211de7d830a324bb5e95a72b1a7b982afb5d926cf853c610ce48 79.9 MiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
docker.io/rancher/mirrored-pause:3.6 application/vnd.oci.image.manifest.v1+json sha256:16974531848218d24822bf606be022d030ab8c9b05b2ecf11076c4c1c6885c95 669.8 KiB linux/amd64 io.cattle.k3s.pinned=pinned,io.cri-containerd.image=managed,io.cri-containerd.pinned=pinned
sha256:0392ee038903218dcd9765e0a0970ea34d07da25da8ccefb17b254be1355d6c
application/vnd.oci.image.manifest.v1+json
sha256:7a6fa2b0e04fb718f81c554cd569ae76dca021205551c20e6falf2f0a2106870 71.7 MiB linux/amd64 io.cri-containerd.image=managed
sha256:180d1ef27ac954c0f033479ec03994076799a9990cd4b9bc4a4b2492982e539a
application/vnd.oci.image.manifest.v1+json
sha256:ef2a7527aea909ee1299efccbad971b7f86e8682311d07970972f57ac6f3e9b5 206.0 MiB linux/amd64 io.cri-containerd.image=managed
sha256:1e6d89f15f8d1cb1ca97733c7053bcd1df5aa9959012e7f985ee0dd9d6690e90
application/vnd.oci.image.manifest.v1+json
sha256:799d5ce61ed32420a80afd9fb027e8932b210b82dffc4cbaa6acbaefe5c74c77 116.0 MiB linux/amd64 io.cri-containerd.image=managed
sha256:2d61ae04c2b80e5421b176b2bb550daab8c4e78e007ea65c9695f55102a3495e
application/vnd.oci.image.manifest.v1+json
sha256:21a5bc93624c80bf8022b7bfad82a3242fa380bb9869300fe5ab4326311ae0ff 4.3 MiB linux/amd64 io.cri-containerd.image=managed
sha256:3a1e150bf4c5610ac288b87c5e42e9ddc514e72b7ecdc43bde4776bda41edbc9
application/vnd.oci.image.manifest.v1+json
sha256:769d5df52751c361a149921905c7dff87f5c3033e8118883e7368d60ffaabc91 214.5 MiB linux/amd64 io.cri-containerd.image=managed

```

```

sha256:6270bb605e12e581514ada5fd5b3216f727db55dc87d5889c790e4c760683fee
application/vnd.oci.image.manifest.v1+json
sha256:16974531848218d24822bf606be022d030ab8c9b05b2ecf11076c4c1c6885c95 669.8 KiB linux/amd64 io.cri-
containerd.image=managed
sha256:8309ed19e06b99d27ea8ade9635fc3aaec0dfaf906fcf71706a679ea444df01f
application/vnd.oci.image.manifest.v1+json
sha256:74de96b9971dd4ea1a835687b80da0619820b6eb41f6f8f8032bdcfe05e791db 57.9 MiB linux/amd64 io.cri-
containerd.image=managed
sha256:b9e1e3849e07022817ebc1612858382f0c0b91d00e4dcd2996adc1df6ced26e9
application/vnd.oci.image.manifest.v1+json
sha256:4557c70a82b0211de7d830a324bb5e95a72b1a7b982afb5d926cf853c610ce48 79.9 MiB linux/amd64 io.cri-
containerd.image=managed
sha256:f7415d0003cb62ded390ed491fc842ee821878a04cc137196c21c1050101dd5e
application/vnd.oci.image.manifest.v1+json
sha256:89a128a64d0f07434abddf496ce78d20208accc263414668bb6d7a978c220226 12.1 MiB linux/amd64 io.cri-
containerd.image=managed
OcNOS#

```

```

OcNOS#exec-shell kubectl apply -f netconfd-pod.yaml
pod/netconfd-pod created

```

```
---
```

```

OcNOS#exec-shell kubectl get pods
NAME          READY   STATUS    RESTARTS   AGE
netconfd-pod  1/1     Running   0           13s
OcNOS#

```

```
---
```

```

OcNOS#exec-shell kubectl apply -f netconfd-svc-lb.yaml
service/netconfd-svc-lb created

```

```
---
```

```

OcNOS#exec-shell kubectl get svc
NAME          TYPE          CLUSTER-IP    EXTERNAL-IP    PORT(S)          AGE
kubernetes    ClusterIP     203.0.114.1   <none>         443/TCP          39m
netconfd-svc-lb LoadBalancer 203.0.114.215 240.0.0.2      10830:32275/TCP 9s

```

```
---
```

```

OcNOS#exec-shell kubectl get nodes
NAME          STATUS    ROLES          AGE   VERSION
ocnos-node    Ready     control-plane,master 42m   v1.33.5+k3s1

```

```
---
```

```

OcNOS#exec-shell kubectl get namespace

```

```

NAME          STATUS    AGE
containers-ns  Active    42m
default        Active    43m
kube-node-lease Active    43m
kube-public    Active    43m
kube-system    Active    43m

```

```
---
```

```

OcNOS#exec-shell kubectl get pods -o wide

```

```

NAME          READY   STATUS    RESTARTS   AGE   IP             NODE          NOMINATED
NODE    READINESS GATES
netconfd-pod  1/1     Running   0           6m40s  203.1.114.6   ocnos-
node    <none>   <none>

```

```
---
```

```

OcNOS#exec-shell kubectl get svc -o wide

```

```

NAME          TYPE          CLUSTER-IP    EXTERNAL-IP    PORT(S)          AGE   SELECTOR
kubernetes    ClusterIP     203.0.114.1   <none>         443/TCP          46m   <none>
netconfd-svc-lb LoadBalancer 203.0.114.215 240.0.0.2      10830:32275/TCP 6m22s app=netconfd

```

```
---
```

```

OcNOS#exec-shell kubectl get nodes -o wide

```

```

NAME          STATUS    ROLES          AGE   VERSION          INTERNAL-IP    EXTERNAL-IP    OS-
IMAGE          KERNEL-VERSION CONTAINER-RUNTIME
ocnos-node    Ready     control-plane,master 52m   v1.33.5+k3s1     240.0.0.2      240.0.0.2      Debian
GNU/Linux 12 (bookworm) 6.1.148-g410b414d0 containerd://2.1.4-k3s1

```

```

OcNOS#

```

```
---
```

```

OcNOS#exec-shell kubectl get namespace -o wide

```

```

NAME          STATUS    AGE
containers-ns  Active    52m
default        Active    52m

```

```
kube-node-lease    Active    52m
kube-public        Active    52m
kube-system        Active    52m
OcNOS#
```

K3s Containers Commands

This chapter is a reference for the K3s feature commands.

feature k3s

Use this command to enable the K3s feature.

Use the `no` parameter to disable the K3s feature.

Command Syntax

```
feature k3s
no feature k3s
```

Parameter

None

Default

None

Command Mode

Global Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

There are two methods to enable the k3s feature:

Type 1: In OcNOS configuration mode and commits it to the running config.

```
OcNOS(config)#feature k3s
OcNOS(config)#commit

OcNOS(config)#no feature k3s
OcNOS(config)#commit
```

Type 2: Linux shell command from OcNOS CLI to start the K3s service using `systemctl`.

```
OcNOS#exec-shell systemctl start k3s.service
```

feature docker

Use this command to enable the docker feature.

Use the `no` parameter to disable the docker feature.

Command Syntax

```
feature docker
no feature docker
```

Parameter

None

Default

None

Command Mode

Global Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
OcNOS(config)#feature docker
OcNOS(config)#commit

OcNOS(config)#no feature docker
OcNOS(config)#commit
```

show feature k3s

Use this command to display the current status of the k3s feature.

Command Syntax

```
show feature k3s
```

Parameter

None

Default

None

Command Mode

Global Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
OcNOS(config)#show feature k3s
k3s                enabled
```

show feature docker

Use this command to display the current status of the docker feature.

Command Syntax

```
show feature docker
```

Parameter

None

Default

None

Command Mode

Global Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
OcNOS(config)#show feature docker
docker                enabled
```

Glossary

Key Terms/Acronym	Description
K3S	K3s is a lightweight, certified Kubernetes distribution designed for resource-constrained environments like the edge and IoT.
Network Address Translation (NAT)	It is a method used by routers and firewalls to modify the network address information (specifically IP addresses and port numbers) in the header of IP packets as they pass through a traffic-routing device.
Container Network Interface (CNI)	It is a set of rules or specifications that defines the standard way container platforms connect to the network, making portable and ensuring to have their own unique IP address and proper network access.
Containers	It is a self contained package of software that bundles an application and its requirements. They share the host OS to keep running the applications separately in isolation from other containers and host.
Orchestration	It is a automated configuration manages and coordinates entire lifecycle of containers across machines or clusters.

Key Terms/Acronym	Description
Pods	These are the smallest deployable units represent a single instance of a running process in the cluster. one or more containers within pod runs together, share resources are treated as a single application unit.

NetConf Access Control Model User Guide

Overview

The NETCONF Access Control Model (NACM) provides a standardized framework for managing user access and permissions within the NETCONF environment. It defines how access to configuration and operational data is controlled, ensuring that only authorized users or groups can view, modify, or execute specific operations on the device.

NACM enables administrators to define fine-grained permissions for different users through both rule-based and group-based access control. It governs which RPCs and configuration data can be viewed or modified. It supports multiple rule types that applies to modules, protocol operations, data nodes, and notifications to offer flexible and precise policy enforcement.

By applying NACM, network devices can be managed more securely and consistently. It helps prevent unauthorized configuration changes, ensures compliance with organizational policies, and aligns with Internet Engineering Task Force(IETF) security standards for NETCONF protocol.

Feature Characteristics

- NACM manage the roles specific permission access to read, write, and execute operation in network devices.
- **User and Group Management:**
 - **ROOT:** The Root user is a super user with unrestricted access.
 - **admin/ocnos User:** admin/ocnos users belongs to **PRIV1 group**, which has all the permission. They can create group, add users to the group, and configure NACM rules for those groups.
- **Restricted Operations:** Only PRIV1 group users (admin/ocnos) and root user can perform `copy-config` and `delete-config` operations.
- **Configuration Persistence:** To ensure NACM configurations are retained across reboots, admin and ocnos users must perform the `<copy-config>` operation with `source=running` and `target=startup`.
- **Recovery:**
 - The super user **root** with unrestricted access and is not bound by NACM rules. Any NetConf session established with the root user is considered a recovery session. During recovery, the root user can **create, delete, or update** one or more NACM rules to bring the device back to a stable state.
 - The admin/ocnos users belong to the PRIV1 group, which has full access permissions through a NACM rule that grants complete privileges to this group. During recovery, the admin/ocnos user can also create, delete, or update one or more NACM rules to restore system stability, provided the PRIV1 rule itself is not deleted.
 - The Root, admin and ocnos users can execute the `delete-configtarget=startup` operation to restore the startup configuration to its default state during recovery scenarios.
- Implemented as a YANG module (`ietf-netconf-acm`) and works with NetConf servers to dynamically enforce access controls.
- **Rule-Based Access Control** access is controlled based on the following rule components:
 - **Target:** The rule applies to which specifies the data nodes, RPCs or notifications.
 - **Action:** Defines the access to permit or deny.

- User/Group: Identifies the entity to which the rule applies.

Role-to-Permission Mapping in NACM

Role/User	Group	Permissions
Root	None	Full unrestricted access to all NetConf operations and configurations.
admin/ocnos	PRIV1	Full access including privileged operations like <code>copy-config</code> , <code>delete-config</code> . admin and ocnos users belongs to PRIV1 group.
Other Users	Custom	Access defined by group-specific NACM rules (for example: read-only, limited RPCs).

Configuration

Initial NACM Configuration

When the NetConf server starts, it is equipped with a default NACM configuration if no prior existing configuration. The initial configuration is as follows:

- By default, the NACM feature is disabled. To enable this feature as per the requirements, only the Root user, admin or ocnos user can send an `edit-config` request to set the configuration `/nacm/enable-nacm=true`.
- Once enabled, if no prior configuration exists under the `/nacm` subtree, the server will deny **read**, **write**, and **execute** access to all operations and data, except for users in the PRIV1 group (**admin/ocnos**) and the Super user **root**.

```
<?xml version="1.0" encoding="UTF-8"?>
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
    <enable-nacm>false</enable-nacm>
    <read-default>deny</read-default>
    <write-default>deny</write-default>
    <exec-default>deny</exec-default>
    <enable-external-groups>false</enable-external-groups>
    <groups>
      <group>
        <name>PRIV1</name>
        <user-name>admin</user-name>
        <user-name>ocnos</user-name>
      </group>
    </groups>
    <rule-list>
      <name>admin-rules</name>
      <group>PRIV1</group>
      <rule>
        <name>permit-all</name>
        <action>permit</action>
        <comment>Permit everything for PRIV1 group</comment>
      </rule>
    </rule-list>
  </nacm>
</config>
```

NetConf NACM vs CLI Role-Based Access

Category	NETCONF NACM	CLI
Role Management	Manages roles and permissions	Manages roles and permissions

	independently.	independently.
Role Definitions	Roles defined via YANG model (ietf-netconf-acm).	Use its own role structure
Access Control Enforcement	Applies rules to NetConf operations and data nodes.	Applies rules to CLI commands.
Authentication Methods	Uses NetConf-specific authentication mechanisms.	CLI uses separate authentication mechanisms.

NACM Rules types

- **Module Rules:** These govern access control to all definitions within the YANG module.

Example: Allow read access to the ipi-interface module.

- **Operation Rules (RPC rule):** These rules restrict access to specific protocol RPC operations or YANG actions. They are defined by the module and the RPC identifier.

Example: Deny access to `<edit-config>` for non-admin users.

- **Notification Rules:** These manages access to specific notification event type, scoped by module and notification name.

Example: Allow access to “interface-link-state-change-notification” notification for operators.

- **Data Rules:** These rules provide fine-grained access control over configuration and operational data via XPath expressions.

Example: Grant read-only access to `/interfaces/interface[name='eth0']`.

Benefits

- Fine-grained control over configuration and operational data.
- Prevents unauthorized changes or sensitive data exposure.

Prerequisites

- The NetConf client should include the NACM capability `urn:ietf:params:xml:ns:yang:ietf-netconf-acm` in its `<hello>` message to use the NACM feature.
- NACM must be enabled in the server configuration.
- User accounts and their corresponding group memberships should be configured before applying NACM rules, as access control is based on user and group identities.

Common NACM Rule Fields

Field	Description
<code>rule-name</code>	Unique name of the NACM rule. This is a required identifier and must be unique within the list of rules.
<code>module-name</code>	The name of the YANG module where the target node, RPC, action, or notification resides (e.g., <code>ietf-interfaces</code> , <code>ietf-netconf</code>). * in <code>module-name</code> allows rules to apply for all modules.

	This is default value.
access-operations	<p>Specifies the NetConf operation types this rule applies to. Can be a space-separated list of any combination of:</p> <ul style="list-style-type: none"> • create – Create a node • read – Read data (get, get-config, notification) • update – Modify existing config • delete – Remove a node • exec – Execute an RPC or action <p>Special value:</p> <ul style="list-style-type: none"> • * – Match all operations.
action	<p>The decision NACM will take when this rule matches.</p> <ul style="list-style-type: none"> • permit – Allow access. • deny – Deny access.
path	XPath expression identifying the data node(s) this rule applies to. Optional.
rpc-name	<p>Name of the RPC or action (e.g., edit-config, get, or custom RPCs). Used only for exec operations. Optional.</p> <p>* – Match all rpc-names.</p>
notification-name	<p>Name of the notification node this rule applies to. Optional.</p> <p>* – Match all notification-names.</p>
comment	Optional comment/description for human readability.
user-name	Username to which this rule applies.
group-name	NACM group to which this rule applies.



Note: Only one of `path`, `rpc-name`, or `notification-name` can be specified in a rule. They are mutually exclusive and depend on the rule type:

- Use `path` for data nodes
- Use `rpc-name` for RPCs or actions
- Use `notification-name` for notifications

Valid Path Notes for NETCONF NACM <path> Rules

When defining NACM <path> rules in NetConf, it is critical to use fully qualified and absolute **XPath expressions** that accurately represent the data model defined in your YANG modules. Follow the best practices below:

General Guidelines

- The <path> must be an absolute XPath, i.e., it must start with /.
- Use fully qualified XPath with correct namespace prefixes and declarations.
- Ensure that all prefixes and namespace URIs match those defined in the corresponding YANG modules.

- Always include prefixes for all keywords including keys inside predicates.

Example 1: Path from Single YANG Module

YANG Module: openconfig-platform

```
module openconfig-platform {
  namespace "http://openconfig.net/yang/platform";
  prefix "oc-platform";
}
```

Valid NACM Rule:

```
<nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
  <rule-list>
    <name>grp_several-rule-list</name>
    <group>grp_several</group>
    <rule>
      <name>grp_several-rule-1</name>
      <path xmlns:oc-platform="http://openconfig.net/yang/platform">
        /oc-platform:components/oc-platform:component
      </path>
      <access-operations>read</access-operations>
      <action>permit</action>
      <comment>grp_several-rule-1-addition</comment>
    </rule>
  </rule-list>
</nacm>
```

Example 2: Path with Augmentation Across Multiple Modules

YANG Modules:

```
module openconfig-terminal-device {
  namespace "http://openconfig.net/yang/terminal-device";
  prefix "oc-opt-term";
}
```

Valid NACM Rule:

```
<nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
  <rule-list>
    <name>grp_several-rule-list</name>
    <group>grp_several</group>
    <rule>
      <name>grp_several-rule-2</name>
      <path xmlns:oc-platform="http://openconfig.net/yang/platform"
        xmlns:oc-opt-term="http://openconfig.net/yang/terminal-device">
        /oc-platform:components/oc-platform:component/oc-opt-term:optical-channel/oc-opt-term:config/oc-opt-term:frequency
      </path>
      <access-operations>read</access-operations>
      <action>permit</action>
      <comment>grp_several-rule-2-addition</comment>
    </rule>
  </rule-list>
</nacm>
```

Example 3: Path with Keys in Predicates (Prefix Required)

Keys must include the appropriate prefix, to describe the function of the key.

Valid XPath Example:

```
<path xmlns:oc-platform="http://openconfig.net/yang/platform"
  xmlns:oc-opt-term="http://openconfig.net/yang/terminal-device">
  /oc-platform:components/oc-platform:component[oc-platform:name='OCH-0/1']/oc-opt-term:optical-channel
</path>
```

Creating NetConf RPC for NACM

RPC Configurations for NACM

Edit-config RPC for enabling NACM

The `edit-config` RPC is used to enable NACM by updating the relevant configuration in the YANG module.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <enable-nacm>true</enable-nacm>
      </nacm>
    </config>
  </edit-config>
</rpc>
```

Edit-config RPCs for group

An RPC is used to **create the user group "PRIV2"** and add the user test to it in the YANG module.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <groups>
          <group>
            <name>PRIV2</name>
            <user-name>test</user-name>
          </group>
        </groups>
      </nacm>
    </config>
  </edit-config>
</rpc>
```

NACM RPCs for Module Rule

An `edit-config` RPC is used to **permit read access** to the **ipi-interface** module for the user group **PRIV2** by configuring rule in the YANG module.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>PermitReadInterfaces</name>
            <module-name>ipi-interface</module-name>
            <access-operations>read</access-operations>
            <action>permit</action>
            <comment>Permit Read Access on "ipi-interface" Module for Group "PRIV2"</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>
```

```

    </nacm>
  </config>
</edit-config>
</rpc>

```

An edit-config RPC is used to **deny read access** to the **ipi-interface** module for the user group **PRIV2** by configuring deny rule in the YANG module.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>DenyReadInterfaces</name>
            <module-name>ipi-interface</module-name>
            <access-operations>read</access-operations>
            <action>deny</action>
            <comment>Deny Read Access on "ipi-interface" Module for Group "PRIV2"</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

An edit-config RPC is used to **permit read access to all modules** for the user group **PRIV2** by configuring YANG module rule.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>PermitReadAllModules</name>
            <module-name>*</module-name>
            <access-operations>read</access-operations>
            <action>permit</action>
            <comment>Permit Read Access on all Modules for Group "PRIV2"</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

An edit-config RPC is used to **permit all operations** on the **ipi-interface** module for the user group **PRIV2** by configuring rule in the YANG module.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>

```

```

<nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
  <rule-list>
    <name>PRIV2-rules</name>
    <group>PRIV2</group>
    <rule>
      <name>PermitAllOperationsInterfaces</name>
      <module-name>ipi-interface</module-name>
      <access-operations>*</access-operations>
      <action>permit</action>
      <comment>Permit all operations on "ipi-interface" Module for Group "PRIV2"</comment>
    </rule>
  </rule-list>
</nacm>
</config>
</edit-config>
</rpc>

```

An edit-config RPC is used to **permit both read and exec access** to the **ipi-interface** module for the user group **PRIV2** by defining a rule in the YANG module.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>PermitReadExecInterfaces</name>
            <module-name>ipi-interface</module-name>
            <access-operations>read exec</access-operations>
            <action>permit</action>
            <comment>Permit Read and exec Access on "ipi-interface" Module for Group
"PRIV2"</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

NACM RPCs for RPC Rule

An edit-config RPC is used to **permit the get-config** operation for the user group **PRIV2** by adding an exec permission rule in the YANG module.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>permit-get-config-rpc</name>
            <rpc-name>get-config</rpc-name>
            <access-operations>exec</access-operations>
            <action>permit</action>
            <comment>Permit get-config rpc for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

```

    </nacm>
  </config>
</edit-config>
</rpc>

```

An edit-config RPC is used to **deny** the get-config operation for the user group **PRIV2** by configuring a rule in the YANG module with access-operations set to exec and action set to deny.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>deny-get-config-rpc</name>
            <rpc-name>deny-config</rpc-name>
            <access-operations>exec</access-operations>
            <action>deny</action>
            <comment>Deny get-config rpc for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

An edit-config RPC is used to **permit all RPC operations for the user group PRIV2** by configuring a wildcard exec rule in the YANG module.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>permit-all-rpcs</name>
            <rpc-name>*</rpc-name>
            <access-operations>exec</access-operations>
            <action>permit</action>
            <comment>Permit all rpcs for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

NACM RPCs for Notification Rule

An edit-config RPC is used to **permit** the interface-link-state-change-notification notification for the user group **PRIV2** by adding a rule in the YANG module with access-operations set to read and action set to permit for notification-name interface-link-state-change-notification.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>

```



```

    <candidate/>
  </target>
</config>
<rpc xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
  <rule-list>
    <name>PRIV2-rules</name>
    <group>PRIV2</group>
    <rule>
      <name>permit-interface-link-state-change-notification</name>
      <notification-name>interface-link-state-change-notification</notification-name>
      <access-operations>read</access-operations>
      <action>permit</action>
      <comment>Permit notification interface-link-state-change-notification for PRIV2
group</comment>
    </rule>
  </rule-list>
</nacm>
</config>
</edit-config>
</rpc>

```

An edit-config RPC is used to **deny** the interface-link-state-change-notification notification for the user group **PRIV2** by configuring a rule in the YANG module with access-operations **set** to read and action **set** to deny.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>deny-interface-link-state-change-notification</name>
            <notification-name>interface-link-state-change-notification</notification-name>
            <access-operations>read</access-operations>
            <action>deny</action>
            <comment>Deny notification interface-link-state-change-notification for PRIV2
group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

An edit-config RPC is used to **permit all notifications** for the user group **PRIV2** by adding a wildcard rule in the YANG module with access-operations **set** to read and action **set** to permit.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>permit-all-notifications</name>
            <notification-name>*</notification-name>
            <access-operations>read</access-operations>
            <action>permit</action>

```

```

        <comment>Permit all notifications for PRIV2 group</comment>
      </rule>
    </rule-list>
  </nacm>
</config>
</edit-config>
</rpc>

```

NACM RPCs for Data Rule

An edit-config RPC is used to **permit all operations** on the `/interfaces` data path for the user group **PRIV2** by configuring a rule in the YANG module.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>permit-xpath-interfaces-all-operations</name>
            <path xmlns:ipi-interface="http://www.ipinfusion.com/yang/ocnos/ipi-interface">/ipi-
interface:interfaces</path>
            <access-operations>*</access-operations>
            <action>permit</action>
            <comment>Permit all operations for xpath interfaces for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```



Note: This rule is applicable for the data path `/interfaces` and all of its descendants.

An edit-config RPC is used to **deny edit operations** on the `/interfaces` data path for the user group **PRIV2** by configuring a rule in the YANG module with `access-operations` set to create, update, and delete and action set to deny.

```

<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>deny-xpath-interfaces-edit-operations</name>
            <path xmlns:ipi-interface="http://www.ipinfusion.com/yang/ocnos/ipi-interface">/ipi-
interface:interfaces</path>
            <access-operations>create update delete</access-operations>
            <action>deny</action>
            <comment>Deny edit operations for xpath interfaces for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>

```

```
</edit-config>
</rpc>
```



Note: This rule is applicable to the data path `/interfaces` and all of its descendants.

The following rules define fine-grained access control for the user group **PRIV2** using XPath expressions, specifically targeting interface-level permissions within the `ipi-interface` YANG module.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>Deny-xpath-interfaces-eth0-edit-operations</name>
            <path xmlns:ipi-interface="http://www.ipinfusion.com/yang/ocnos/ipi-interface">/ipi-
interface:interfaces/ipi-interface:interface[ipi-interface:name="eth0"]</path>
            <access-operations>create update delete</access-operations>
            <action>deny</action>
            <comment>deny edit operations for xpath interface eth0 for PRIV2 group</comment>
          </rule>
          <rule>
            <name>permit-xpath-interfaces-edit-operations</name>
            <path xmlns:ipi-interface="http://www.ipinfusion.com/yang/ocnos/ipi-interface">/ipi-
interface:interfaces</path>
            <access-operations>create update delete</access-operations>
            <action>permit</action>
            <comment>Permit edit operations for xpath interfaces for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>
```

Purpose of Rules

This example demonstrates how to restrict edit operations (`create`, `update`, `delete`) only on interface `eth0` while allowing those operations on all other interfaces for users in NACM group `PRIV2`.

Rules Breakdown

- The **first rule** denies edit operations specifically on the interface node where `name = "eth0"`.
- The **second rule** permits edit operations on all interfaces under the **XPath** `/ipi-interface:interfaces`.

Rule Order

- NACM rules are evaluated in the order they appear.
- In this example, the deny rule comes first, so when the server evaluates access for `eth0`, it finds a match and denies the operation before reaching the permit rule.
- If the permit rule is placed before the deny rule, the server would match it first (because `/ipi-interface:interfaces` includes `eth0` as a descendant) and would therefore incorrectly allow edit operations on `eth0`.

Best Practice

- Always define more specific rules (e.g., for a particular interface) before more general ones (e.g., for the entire interfaces list).
- This ensures that exceptions are enforced before broader access is granted.

Rule Insertion Control

To insert a NACM rule at a specific position (e.g., before or after another rule), you can use the `yang:insert` attribute as explained in [Ordered Rule Management with Yang:Insert](#) of this guide.

Edit-config RPC to **Permit All Operations** on `/interfaces/interface`, Except Interfaces start with "eth", for user group **PRIV2**.

```
<rpc xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <edit-config>
    <target>
      <candidate/>
    </target>
    <config>
      <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
        <rule-list>
          <name>PRIV2-rules</name>
          <group>PRIV2</group>
          <rule>
            <name>Deny-xpath-interfaces-eth-edit-operations</name>
            <path xmlns:ipi-interface="http://www.ipinfusion.com/yang/ocnos/ipi-interface">/ipi-
interface:interfaces/ipi-interface:interface[ipi-interface:name='eth.*']</path>
            <access-operations>create update delete</access-operations>
            <action>deny</action>
            <comment>deny edit operations for xpath interface strat with eth for PRIV2
group</comment>
          </rule>
          <rule>
            <name>permit-xpath-interfaces-edit-operations</name>
            <path xmlns:ipi-interface="http://www.ipinfusion.com/yang/ocnos/ipi-interface">/ipi-
interface:interfaces</path>
            <access-operations>create update delete</access-operations>
            <action>permit</action>
            <comment>Permit edit operations for xpath interfaces for PRIV2 group</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
</rpc>
```

Enforce Access Control for OpenConfig Data Models

- NetConf Access Control Model (NACM) enforce access control validation for OpenConfig YANG data models when OpenConfig translation is enabled.
- This ensures that only authorized users can access, modify, or execute configuration and operational data from OpenConfig models in a network device or controller.
- Admin, ocnos, or root users can configure NACM rules for OpenConfig data models when OpenConfig translation is enabled, as explained in the previous section.

Ordered Rule Management with Yang:Insert

- NACM's `rule-list` and `rule` elements are defined as ordered-by user, meaning that administrators must have the ability to define rule precedence explicitly.
- Below are examples demonstrating how to insert rules using `yang:insert`:

Initial NACM Rule Configuration:

```
<config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
    <rule-list>
      <name>admin-rules</name>
      <rule>
        <name>existing-rule</name>
        <action>permit</action>
        <comment>Existing rule in NACM</comment>
      </rule>
    </rule-list>
  </nacm>
</config>
```

Adding a New Rule After an Existing Rule:

```
<edit-config>
  <config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
    <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
      <rule-list>
        <name>admin-rules</name>
        <rule
          xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm"
          xmlns:yang="urn:ietf:params:xml:ns:yang:1"
          yang:insert="after"
          yang:key="[name='existing-rule']">
            <name>new-rule</name>
            <action>deny</action>
            <comment>New rule inserted after existing-rule</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
```

Adding a Rule Before an Existing Rule:

```
<edit-config>
  <config xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
    <nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
      <rule-list>
        <name>admin-rules</name>
        <rule
          xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm"
          xmlns:yang="urn:ietf:params:xml:ns:yang:1"
          yang:insert="before"
          yang:key="[name='existing-rule']">
            <name>high-priority-rule</name>
            <action>deny</action>
            <comment>Inserted before existing-rule</comment>
          </rule>
        </rule-list>
      </nacm>
    </config>
  </edit-config>
```

Adding a Rule at first:

```
<nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
  <rule-list>
    <name>admin-rules</name>
    <rule xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm"
      xmlns:yang="urn:ietf:params:xml:ns:yang:1"
      yang:insert="first">
      <name>allow-get-config</name>
      <rpc-name>get-config</rpc-name>
      <access-operations>*</access-operations>
      <action>permit</action>
    </rule>
  </rule-list>
</nacm>
```

```

    <comment>Allow get-config rpc for PRIV1 group</comment>
  </rule>
</rule-list>
</nacm>

```

Adding a Rule at last:

```

<nacm xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm">
  <rule-list>
    <name>admin-rules</name>
    <rule xmlns="urn:ietf:params:xml:ns:yang:ietf-netconf-acm"
      xmlns:yang="urn:ietf:params:xml:ns:yang:1"
      yang:insert="last">
      <name>allow-get-config</name>
      <rpc-name>get-config</rpc-name>
      <access-operations>*</access-operations>
      <action>permit</action>
      <comment>Allow get-config rpc for PRIV1 group</comment>
    </rule>
  </rule-list>
</nacm>

```

Using NETCONF to Manage NACM Rules

- Apply rules using the <edit-config> RPC.
- Validate rule application using <get-config>.
- Verify rule enforcement by performing specific operations.
- Modify or delete rules using <edit-config> with the appropriate XPath.

Troubleshooting

- Ensure the rule list name and paths are correct.
- Confirm users are part of the correct groups. Check netconfd logs for errors.
- Enable NETCONF debug logs as given below and check

```
yangcli ocnos@0> set-log-level log-level=debug4
```



Note:

- To check if the rule is applied or not - Use <get-config> to retrieve the NACM configuration.
- To delete a rule - Use <edit-config> with the operation delete on the target rule path.

Limitations

- NACM relies on NetConf transport layer for user authentication.
- User-to-group mapping is dependent on admin.
- Performance impacts on large configurations, NACM rule evaluation can add processing overhead, especially for fine-grained data node checks.

Glossary

The following provides definitions for key terms used throughout this document.

Access Control	A security feature provided by the server that allows an administrator to restrict access to a subset of all protocol operations and data, based on various criteria.
NETCONF Access Control Model (NACM)	A model used to configure and monitor the access control procedures desired by the administrator to enforce a particular access control policy.
YANG Module	YANG (Yet Another Next Generation) is a data modeling language standardized by the IETF (RFC 7950). It is a structured, machine-readable file that defines the data model used by network management protocols notably NETCONF.
Remote Procedure Calls (RPC) rule	In the context of NACM (Network Configuration Access Control Model) it is an access control rule that determines whether a user is allowed or denied permission to run specific RPCs or actions defined in YANG modules.
NETCONF	Network Configuration Protocol
RPC	Remote Procedure Call
TLS	Transport Layer Security
SSH	Secure Shell
PRIV1	Highest Privilege user group (admin and ocnos users)

sFlow - Sample Packet Monitoring for Multiple Interfaces

Overview

This chapter provides the steps for configuring Sampled Flow (sFlow).

sFlow is the standard for monitoring high-speed switches and routes in a network. It collects sample traffic from high-speed network devices to calculate its performance statistics. The sFlow system consists of an sFlow Agent which is embedded in a switch or router and an sFlow Collector.

The sFlow agent samples packets on both ingress and egress directions as well as polling traffic statistics for the device it is monitoring. The packet sampling is performed by the switching/routing device at wire speed. The sFlow agent forwards the sampled traffic statistics in sFlow Packet Data Units (PDUs) as well as sampled packets to an sFlow collector for analysis.



Note: sFlow egress sampling for multicast, broadcast, or unknown unicast packets is not supported.

The sFlow agent uses the following forms of sampling:

- Sampling packets: samples one packet out of a defined sampling rate. This sampling is done by hardware at wire speed.
- Sampling counters: polls interface statistics such as generic and Ethernet counters at a defined interval.

The sFlow feature collects sampled traffic data and counters from configured interfaces. The collected data is sent to all collectors (by default) using the sFlow protocol. For more information, refer to [RFC 3176](#).

This functionality support multiple collectors for interfaces simultaneously.

Features Characteristics

- Supports maximum of five concurrent sFlow collectors on the system.
- Uses a specific user defined VRF interface for each collector. If not specified, the management VRF is used.
- Sends the collected sFlow samples on each interface to all configured collectors on the system.
- Has the ability to disable the sending of sFlow samples from an interface to specified collectors.
- sFlow sampling monitoring can be enabled globally across all interfaces with a single command.
- The sFlow feature is supported on both physical interfaces and LAG (Link Aggregation Group) interfaces. When sampling is configured on a LAG interface, it is automatically applied to all member ports within that LAG.
- When sFlow sampling is in-progress on high rate, CPU usage spike messages from Chassis monitoring module (cmmd) is expected.
- The Qumran 1 platform is equipped to handle a total of 9 unique sampling rates. Ingress and egress sampling rate is counted separately.
- The Qumran 2 platform is equipped to handle a total of 15 unique sampling rates.
 - For egress, maximum 7 unique sampling rates can be created.
 - If egress sampling is not used, a total of 15 unique ingress sampling rates can be configured.
 - Total ingress sampling = 15 - number of egress sampling rates.

Benefits

The sFlow with multiple collectors provides the capability to do multiple service analysis simultaneous in a network.

Tracks network utilization, bandwidth usage, and performance metrics across interfaces.

Analyzes traffic flows to understand application usage, user behavior, or device interactions.

Prerequisites

Make sure to enable the required interface with sflow feature and an agent IP address.

```
feature sflow
sflow agent-ip 1.2.7.10
interface xe1
  sflow enable
!
```



Configuration

This section provides the configurations required to assign multiple sFlow collectors to all interfaces.

Topology

The following topology illustrates the sFlow multiple collectors connected to multiple interfaces with one sFlow Packet Data Unit (PDU):

Figure 6. sFlow with Multiple Collectors



Perform the following configurations on OcNOS device:

1. Login to Config mode and enable sFlow.

```
#configure terminal
(config)#feature sflow
```

2. Configure the sFlow collector whose IP address must be reachable via the management VRF or VRF default.

```
(config)#sflow collector-id 3 collector 1.2.3.24 port 6345 receiver-time-out 5 max-datagram-size 1560
(config)#sflow collector-id 4 collector 1.2.4.24 port 6346 receiver-time-out 4 max-datagram-size 1570 vrf default
```

3. Configure sFlow attributes including counter poll interval, ingress traffic direction, sampling rate, and maximum header size for sampled packets on interface `xe12` and exit configuration mode.

```
(config)#interface xe12
(config-if)#sflow poll-interval 5
(config-if)#sflow direction ingress
(config-if-sflow)#sampling-rate 1024
(config-if-sflow)#max-header-size 256
(config-if-sflow)#exit
(config-if)#sflow enable
(config-if)#commit
(config-if)#end
```

4. Configure the sampling rate and maximum header size, enabling packet sampling, collector Id on interface xe12 for sFlow egress and exit the configuration mode.

```
(config-if)#sflow direction egress
(config-if-sflow)#sampling-rate 2000
(config-if-sflow)#max-header-size 16
(config-if-sflow)#exit
(config-if)#sflow enable
(config-if)#sflow poll-interval 10
(config-if)#commit
(config-if)#end
```

5. Configure sFlow for other interface xe13.

```
(config-if)#interface xe13
(config-if)#sflow direction ingress
(config-if-sflow)#sampling-rate 2500
(config-if-sflow)#max-header-size 100
(config-if-sflow)#exit
(config-if)#sflow direction egress
(config-if-sflow)#sampling-rate 2000
(config-if-sflow)#max-header-size 16
(config-if-sflow)#exit
OcNOS(config-if)#sflow enable
OcNOS(config-if)#sflow poll-interval 5
```

Show Running Configuration

The following show output displays the sample sFlow configuration details.

```
OcNOS#show running-config sflow
feature sflow
!
sflow agent-ip 1.2.7.10
sflow collector-id 3 collector 1.2.3.24 port 6345 receiver-time-out 5 max-
datagram-size 1560
sflow collector-id 4 collector 1.2.4.24 port 6346 receiver-time-out 4 max-
datagram-size 1570 vrf default
!
interface xe12
sflow enable
sflow direction ingress
sampling-rate 1024
max-header-size 256
exit
sflow direction egress
sampling-rate 2000
max-header-size 16
exit
sflow poll-interval 10
!

interface xe13
sflow enable
sflow direction ingress
sampling-rate 2500
max-header-size 100
exit
sflow direction egress
sampling-rate 2000
max-header-size 16
exit
sflow poll-interval 5
!
```

Validation

The following show output displays the sFlow details:

```
OcNOS#show sflow detail
sFlow Feature: Enabled
sFlow Version: 5
Agent IP      : 1.2.7.10
Collector 3:
  IP: 1.2.3.24      Port: 6345
  VRF                :
  Maximum Datagram Size(bytes): 1560
  Receiver timeout(sec) : 0
Collector 4:
  IP: 1.2.4.24      Port: 6346
  VRF                :
  Maximum Datagram Size(bytes): 1570
  Receiver timeout(sec) : 0
```

sFlow Port Detailed Information:

Interface Polling	Collector	Packet-Sampling Maximum Header	Packet-Sampling	Counter-		
ID	Rate	Count	Interval	Count	Size(bytes)	
Ingress	Egress	Ingress	Egress	(sec)	Ingress	Egress
-----	-----	-----	-----	-----	-----	-----
xe12 256	3 16	1024	2000	3	6	10
xe13 100	4 16	2500	2000	4	7	5
						3

Configuring sFlow with User Defined VRFs

The sFlow feature allows user to configure sample packets using VRF interface.

- Users can sample packets on an interface mapped to a user-defined VRF and send sFlow packets through the same VRF.
- Users can send sampled packets to multiple destinations (collectors) through different VRFs simultaneously.

The following sample configuration demonstrates sFlow using multiple collector-ids with user-defined VRFs:

```
feature sflow

sflow collector-id 3 collector 172.20.1.1 port 6343 receiver-time-out 0 max-datagram-size 200 vrf
sys_mgmt
sflow collector-id 4 collector 192.168.7.2 port 6343 receiver-time-out 1000 max-datagram-size 200 vrf
xe11_vrf
sflow collector-id 5 collector 172.10.1.1 port 65535 receiver-time-out 0 max-datagram-size 200 vrf
xe10_10_vrf
sflow collector-id 2 collector 10.1.1.1 port 1024 receiver-time-out 345 max-datagram-size 400 vrf
xe10_vrf

!

interface xe12
sflow direction ingress
sampling-rate 1029
max-header-size 120
```

```
exit
  sflow direction egress
  sampling-rate 1029
  max-header-size 120
exit

sflow enable

!

interface xe13
sflow direction ingress
  sampling-rate 1048
  max-header-size 140
exit
sflow enable

!

interface xe14
  sflow direction ingress
  sampling-rate 1048
  max-header-size 128
exit
  sflow direction egress
  sampling-rate 1048
  max-header-size 128
exit

sflow enable
sflow poll-interval 20

!

interface xe15
sflow direction ingress
  sampling-rate 1029
  max-header-size 120
exit

sflow enable
```

Validation

The following show output displays the sFlow details associated with multiple VRFs:

```
S9510-30XC-A#show sflow detail

sFlow Feature: Enabled

sFlow Version: 5

Agent IP : 172.16.1.2

Collector 3:

IP: 172.20.1.1 Port: 6343

VRF : sys_mgmt

Maximum Datagram Size(bytes): 200
```

Receiver timeout(sec) : 0

Collector 4:

IP: 192.168.7.2 Port: 6343

VRF : xe11_vrf

Maximum Datagram Size(bytes): 200

Receiver timeout(sec) : 0

Collector 5:

IP: 172.10.1.1 Port: 65535

VRF : xe10_10_vrf

Maximum Datagram Size(bytes): 200

Receiver timeout(sec) : 0

Collector 1:

IP: 192.168.7.2 Port: 65530

VRF :

Maximum Datagram Size(bytes): 200

Receiver timeout(sec) : 0

Collector 2:

IP: 10.1.1.1 Port: 1024

VRF : xe10_vrf

Maximum Datagram Size(bytes): 400

Receiver timeout(sec) : 0

sFlow Port Detailed Information:

Interface Polling	Collector ID	Packet-Sampling Maximum Header	Packet-Sampling Sampling	Packet-Sampling Interval	Counter-	
(s) (bytes)	Rate	Direction	Count	Interval	Count	Size
(sec)		Ingress Ingress	Egress Egress	Ingress Egress		
-----	-----	-----	-----	-----	-----	---
xe12	1	1029	1029	0	0	0
120	120	egress-only				
xe13	5	1048	0	0	0	0
140	0	ingress-only				
xe14	3	1048	1048	0	0	20
128	128	ingress-only				1248
xe15	4	1029	0	0	0	0
120	0	ingress-only				
xe16	2	2048	3020	0	0	0
140	128	egress-only				

Implementation Examples

Example 1

To configure multiple sFlow collectors for multiple interfaces:

```
(config)#feature sflow
(config)#sflow agent-ip 172.16.0.25
(config)#sflow poll-interval 20
(config-sflow)#sflow direction ingress
(config-sflow)#sampling-rate 1024
(config-sflow)#max-header-size 256
(config-sflow)#exit
(config)#sflow direction egress
(config-sflow)#sampling-rate 1024
(config-sflow)#max-header-size 128
(config-sflow)#exit
```

Verify the sFlow collector details:

```
#show sflow detail
sFlow Feature: Enabled
sFlow Version: 5
Agent IP      : 172.16.0.25
Collector 1:
  IP: 8.12.33.201      Port: 6345
  VRF                  :
  Maximum Datagram Size(bytes): 1024
  Receiver timeout(sec) : 0
Collector 2:
  IP: 172.12.33.202    Port: 6343
  VRF                  :
  Maximum Datagram Size(bytes): 1024
  Receiver timeout(sec) : 0
Collector 3:
  IP: 172.12.33.202    Port: 6345
  VRF                  :
  Maximum Datagram Size(bytes): 2048
  Receiver timeout(sec) : 0
Collector 4:
  IP: 172.12.33.202    Port: 7546
  VRF                  :
  Maximum Datagram Size(bytes): 1560
  Receiver timeout(sec) : 0
Collector 5:
  IP: 1.1.3.2          Port: 8998
  VRF                  :
  Maximum Datagram Size(bytes): 1560
  Receiver timeout(sec) : 0
```

sFlow Port Detailed Information:

Interface	Collector	Packet-Sampling	Packet-Sampling	Counter-
Polling	ID	Maximum Header	Sampling	Count
(s)	Rate	Direction	Count	Interval
(bytes)		Ingress	Egress	Count
(sec)		Ingress	Egress	Size
-----	-----	-----	-----	-----
-----	-----	-----	-----	-----

No Interface is enabled for sampling

#

```
#show sflow brief
sFlow Feature: Enabled
Collector 1:
  IP: 8.12.33.201      Port: 6345
Collector 2:
  IP: 172.12.33.202   Port: 6343
Collector 3:
  IP: 172.12.33.202   Port: 6345
Collector 4:
  IP: 172.12.33.202   Port: 7546
Collector 5:
  IP: 1.1.3.2         Port: 8998

sFlow Port Configuration:
Interface  Collector  Status      Sample Rate      Counter-Polling
          ID(s)      Ingress    Egress          Ingress    Egress          Interval(sec)
-----
No Interface is enabled for sampling
#
#
```

Configure multiple sFlow collectors on particular interfaces.



Note: The interface configuration takes precedence over global configuration.

```
(config)#interface xe1
(config-if)#sflow enable
(config-if)#commit
(config-if)#exit
(config)#inter xe2
(config-if)#sflow enable
(config-if)#commit
(config-if)#exit
(config)#interface xe3
(config-if)#sflow enable
(config-if)#commit
(config-if)#exit
```

Verify the sFlow configuration:

```
OcNOS#show sflow
sFlow Feature: Enabled
Collector 1:
  IP: 8.12.33.201      Port: 6345
Collector 2:
  IP: 172.12.33.202   Port: 6343
Collector 3:
  IP: 172.12.33.202   Port: 6345
Collector 4:
  IP: 172.12.33.202   Port: 7546
Collector 5:
  IP: 1.1.3.2         Port: 8998

sFlow Port Configuration:
Interface  Collector  Status      Sample Rate      Counter-Po
lling      ID(s)      Ingress    Egress          Ingress    Egress          lling
          ID(s)      Ingress    Egress          Ingress    Egress          Interval(s
ec)
-----
-----
xe1        1,2,3,4,5  Enabled    Enabled          1024        1024           20
xe20       1,2,3,4,5  Enabled    Enabled          1024        1024           20
OcNOS#
```

```

OcNOS#show sflow detail
sFlow Feature: Enabled
sFlow Version: 5
Agent IP      : 172.16.0.25
Collector 1:
  IP: 8.12.33.201      Port: 6345
  VRF                  :
  Maximum Datagram Size(bytes): 1024
  Receiver timeout(sec) : 0
Collector 2:
  IP: 172.12.33.202    Port: 6343
  VRF                  :
  Maximum Datagram Size(bytes): 1024
  Receiver timeout(sec) : 0
Collector 3:
  IP: 172.12.33.202    Port: 6345
  VRF                  :
  Maximum Datagram Size(bytes): 2048
  Receiver timeout(sec) : 0
Collector 4:
  IP: 172.12.33.202    Port: 7546
  VRF                  :
  Maximum Datagram Size(bytes): 1560
  Receiver timeout(sec) : 0
Collector 5:
  IP: 1.1.3.2          Port: 8998
  VRF                  :
  Maximum Datagram Size(bytes): 1560
  Receiver timeout(sec) : 0

```

sFlow Port Detailed Information:

Interface	Collector	Packet-Sampling	Packet-Sampling	Counter-
Polling	Maximum Header	Sampling	Interval	Count
(s)	ID	Count	Interval	Count
(bytes)	Direction	Ingress	Egress	Size
(sec)	Ingress	Egress	Ingress	Egress
	Ingress	Egress	Ingress	Egress
-----	-----	-----	-----	-----
xe1	1,2,3,4,5	1024	1024	0
256	128	both		0
xe20	1,2,3,4,5	1024	1024	0
256	128	both		0

Example 2

To disable collector(s) on an interface.

```

OcNOS#conf terminal
OcNOS(config)#
OcNOS(config)#interface xe1
OcNOS(config-if)#
OcNOS(config-if)#no sflow collector-id 2
OcNOS(config-if)#no sflow collector-id 4
OcNOS(config)#interface xe20
OcNOS(config-if)#
OcNOS(config-if)#no sflow collector-id 1
OcNOS(config-if)#no sflow collector-id 3
OcNOS(config-if)#no sflow collector-id 5
OcNOS(config-if)#commit
OcNOS(config-if)#end

```

Verify the sFlow collector details:


```
OcNOS#show sflow brief
sFlow Feature: Enabled
Collector 1:
  IP: 8.12.33.201      Port: 6345
Collector 2:
  IP: 172.12.33.202   Port: 6343
Collector 3:
  IP: 172.12.33.202   Port: 6345
Collector 4:
  IP: 172.12.33.202   Port: 7546
Collector 5:
  IP: 1.1.3.2         Port: 8998

sFlow Port Configuration:
Interface  Collector  Status      Sample Rate      Counter-Polling
          ID(s)      Ingress    Egress          Ingress          Egress          Interval (sec)
-----
xe1         1,3,5  Enabled    Enabled         1024             1024             20
xe20        2,4   Enabled    Enabled         1024             1024             20
OcNOS#
OcNOS#
```

Example 3

To remove multiple sFlow collectors.

```
(config)#no sflow collector-id 1 collector 8.12.33.201 port 6345
(config)#no sflow collector-id 3 collector 172.12.33.202 port 6345
(config)#no sflow collector-id 4 collector 172.12.33.202 port 7546
```

Example 4

To verify multiple sFlow collectors by sampling traffic on interfaces in ingress directions.

Ingress direction on global configuration:

```
(sflow-global-config)#sflow direction ingress
(sflow-global-config)#max-header-size 128
(sflow-global-config)#sampling-rate 1024
(sflow-global-config)#commit
(sflow-global-config)#exit
```

Ingress Direction on interface configuration.

```
OcNOS#conf terminal
Enter configuration commands, one per line.  End with CNTL/Z.
OcNOS(config)#interface xe1
OcNOS(config-if)#sflow ingre
OcNOS(config-if)#sflow direction ingress
OcNOS(sflow-if-config)#max-header-size 128
OcNOS(sflow-if-config)#sampling-rate 1024
OcNOS(sflow-if-config)#commit
OcNOS(sflow-if-config)#end
```

Verify the sFlow global and interface ingress configurations.

```
OcNOS#show sflow brief
sFlow Feature: Enabled
Collector 1:
  IP: 172.16.0.100      Port: 6343
Collector 2:
  IP: 172.12.33.202     Port: 9947
Collector 3:
```

```

IP: 192.168.5.73      Port: 6345
Collector 4:
IP: 192.168.5.73      Port: 7546
Collector 5:
IP: 11.0.0.37         Port: 8998

```

sFlow Port Configuration:

Interface	Collector ID(s)	Status		Sample Rate		Counter-Polling Interval (sec)
		Ingress	Egress	Ingress	Egress	
xe1	1,2,3,4,5	Enabled	Disabled	1024	0	0

OcNOS#

OcNOS#show sflow detail

```

sFlow Feature: Enabled
sFlow Version: 5
Agent IP      : 0.0.0.0
Collector 1:
IP: 172.16.0.100      Port: 6343
VRF                :
Maximum Datagram Size(bytes): 1560
Receiver timeout(sec) : 0
Collector 2:
IP: 172.12.33.202     Port: 9947
VRF                :
Maximum Datagram Size(bytes): 1560
Receiver timeout(sec) : 0
Collector 3:
IP: 192.168.5.73      Port: 6345
VRF                :
Maximum Datagram Size(bytes): 2048
Receiver timeout(sec) : 0
Collector 4:
IP: 192.168.5.73      Port: 7546
VRF                :
Maximum Datagram Size(bytes): 1560
Receiver timeout(sec) : 0
Collector 5:
IP: 11.0.0.37         Port: 8998
VRF                :
Maximum Datagram Size(bytes): 1560
Receiver timeout(sec) : 0

```

sFlow Port Detailed Information:

Interface	Collector ID	Packet-Sampling		Packet-Sampling		Counter-	
		Maximum Header	Sampling	Sampling		Count	Size
				Interval			
				Ingress	Egress		
xe1	1,2,3,4,5	1024	0	0	0	0	0
128	0	ingress-only					

OcNOS#

Example 5

To verify multiple sFlow collectors by sampling traffic on interfaces in egress directions.

Egress Direction on Global configuration:

```
(config)#sflow direction egress
```

```
(sflow-global-config)# sampling-rate 2048
(sflow-global-config)# max-header-size 128
(sflow-global-config)# exit
```

Egress Direction on interface configuration.

```
(config)# interface xe1
(config-if)# sflow enable
(config-if)# sflow direction egress
(sflow-if-config)# sampling-rate 1024
(sflow-if-config)# max-header-size 128
(sflow-if-config)# exit
(config-if)# interface xe2
(config-if)# sflow enable
(config-if)# exit
(config)# commit
(config)# exit
```

Verify the sFlow global and interface egress configurations.

```
#show sflow detail
sFlow Feature: Enabled
sFlow Version: 5
Agent IP      : 10.14.111.101
Collector 4:
  IP: 10.0.0.37      Port: 6343
  VRF                : default
  Maximum Datagram Size(bytes): 1560
  Receiver timeout(sec) : 0
Collector 5:
  IP: 11.0.0.37      Port: 7777
  VRF                : default
  Maximum Datagram Size(bytes): 250
  Receiver timeout(sec) : 0
```

sFlow Port Detailed Information:

Interface	Collector ID	Packet-Sampling Maximum Header	Packet-Sampling Count	Packet-Sampling Interval	Counter-Polling Count	Counter-Polling Size
(s)	(bytes)	Rate	Direction	Ingress	Egress	
(sec)			Ingress	Egress	Ingress	Egress
xe1	4,5	0	1024	0	2446	10
0	128	egress-only				
xe2	4,5	0	2048	0	752	0
0	128	egress-only				

```
#show sflow statistics
```

sFlow Port Statistics:

Interface	Collector ID(s)	Packet-Sampling Count	Counter-Polling Count
		Ingress	Egress
xe1	1,2,3,4,5	6629	5798
#			411

Interfaces using sFlow global configuration:

Interface	Packet-Sampling Rate	Polling Interval	Maximum Header Size
-----------	----------------------	------------------	---------------------

	Ingress	Egress		Ingress	Egress
	-----	-----		-----	-----
xe2	no	yes	no	no	yes

Commands

The feature introduces the following configuration command.

- [no sflow collector-id \(page 253\)](#) - When configured with `no`, the `show running-config` output displays the collectors not used by the interface as `no sflow collector-id`.

The following existing commands are modified.

- [sflow collector \(page 254\)](#) - Introduces default values for Port: 6343, Receiver timeout: 0 (no timeout) and Maximum datagram size: 1560.
- [show sflow statistics](#) - Included Collector ID in the output.

For additional information, refer to the [sFlow Commands](#) section.

no sflow collector-id

This command removes the association of a specified sFlow collector from an interface. By default, all sFlow collectors are automatically linked to every interface where sFlow is enabled. With this command, users can control which collectors remain associated with an interface. Because collectors are already in use, removing them may disrupt existing associations and affect ongoing sFlow operations. To re-establish the association, use the [sflow collector-id](#) command.

Use `sflow collector-id` to re-enable the sFlow collector.



Note: This is a negative command. Configures with `no` and displays on `show running-config` as a list of collectors not in use by the interface as `no sflow collector-id`.

Command Syntax

```
no sflow collector-id <1-5>
sflow collector-id <1-5>
```

Parameter

collector-id <1-5>

Specifies the name of the Collector instance identifier.

Default

All sFlow collectors are enabled for all interfaces.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

The following example shows that all sFlow collectors are automatically linked to every interfaces where sFlow is enabled. It also shows that sFlow collector-id 3 and 5 are removed from interface `eth1`.

```
OcNOS#conf t
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#feature sflow
OcNOS(config)#sflow collector 1.1.1.1
OcNOS(config)#sflow collector-id 2 collector 1.1.1.1 port 6344 receiver-time-out 5 max-datagram-size 1256
OcNOS(config)#sflow collector-id 3 collector 1.2.3.4 port 1024 receiver-time-out 60 max-datagram-size 200 vrf default
OcNOS(config)#sflow collector-id 4 collector 1.1.1.1 port 6346 receiver-time-out 1
OcNOS(config)#sflow collector-id 5 collector 2.2.2.2 max-datagram-size 1560
OcNOS(config)#interface eth1
OcNOS(config-if)#sflow enable
OcNOS(config-if)#no sflow collector-id 3
OcNOS(config-if)#no sflow collector-id 5
OcNOS(config-if)#commit
OcNOS(config-if)#end
OcNOS#
```

sflow collector

Use this command to configure the collector details such as the collector IPv4 address, port number, receiver time-out and datagram size.

Use the **no** form of this command to disable the sFlow collector.

Command Syntax

```
sflow (collector-id <1-5>|) collector A.B.C.D (port <1024-65535>|) (receiver-time-out <0-2147483647>|) (max-datagram-size <200-9000>|) (vrf WORD|)  
no sflow collector collector-id <1-5> A.B.C.D port <1024-65535>
```

Parameter

collector-id <1-5>

(Optional) Specifies the name of the Collector instance identifier. If the collector-id is not specified, the ID will be 1.

collector A.B.C.D

Collector IPv4 address. This address must be reachable via the management VRF. <1024-65535>

port <1024-65535>

(Optional) Collector UDP Port number. The default port number is 6343.

receiver-time-out <0-2147483647>

(Optional) Receiver time out value in seconds. Upon timeout, value collector information is removed, stopping any ongoing sampling. The default timeout value is 0 (no timeout).

max-datagram-size <200-9000>

(Optional) Maximum datagram size in bytes that can be sent to the collector. The default value is 1560.

vrf WORD

(Optional) Specifies the User defined VRF to reach the collector. The default used VRF is the management VRF.

Default

Disabled.

Command Mode

Configure mode

Applicability

This command was introduced before OcNOS version 1.3. Introduced the **collector-id** and **vrf** parameters in the OcNOS version 6.5.1. Introduced default values for **port**, **received-time-out**, **max-datagram-size** in OcNOS version 7.0.0

Example

```
#configure terminal
```

```
(config)#sflow collector-id 3 collector 1.2.3.4 port 1024 receiver-time-out 60 max-datagram-size 200  
vrf default  
(config)#no sflow collector
```

```
(config)#interface xel2  
(config-if)#sflow direction ingress  
(config-if-sflow)#sampling-rate 1024
```

```
(config-if-sflow)#max-header-size 256
(config-if-sflow)#exit
(config-if)#sflow enable
(config-if)#sflow poll-interval 10
```

```
OcNOS(config)#sflow collector-id 1 collector 1.1.1.1 port 6343
OcNOS(config)#sflow collector-id 2 collector 1.1.1.1 port 6344 receiver-time-out 5 max-datagram-size 1256
OcNOS(config)#sflow collector-id 3 collector 1.2.3.4 port 1024 receiver-time-out 60 max-datagram-size 200 vrf default
OcNOS(config)#sflow collector-id 4 collector 1.1.1.1 port 6346 receiver-time-out 1
OcNOS(config)#sflow collector-id 5 collector 2.2.2.2 max-datagram-size 1560
```

Troubleshooting

Execute the following commands to check the sFlow configuration at the interface level.

- [show sflow global](#)
- [show sflow brief and detail](#)

Glossary

Key Terms/Acronym	Description
PDU	A unit of data transmitted as a composite by a protocol.
sFlow	Sampled Flow data sFlow (sFlow) is the standard for monitoring high-speed switched and routed networks. The sFlow monitoring system consists of an sFlow Agent which is embedded in a switch or router and an sFlow Collector.

VxLAN OAM for Overlay Networks

OcNOS supports VxLAN Operations, Administration, and Maintenance (OAM) to enhance visibility and fault management for VxLAN overlays in CLOS data center fabric. Using Maintenance End Points (MEPs) at VxLAN Tunnel End Point (VTEPs) and Spines within VxLAN tunnels, operators can perform the following operations to verify connectivity, and isolate faults.

- Ping /Loopback - Verify reachability to a remote VTEP and that the VxLAN tunnel is operational end-to-end.
- Pathtrace - Discover the full forwarding path inside the VxLAN fabric, hop-by-hop
- Continuity checks - Provide continuous, periodic monitoring of VxLAN tunnel health.

The feature supports both static and dynamic VxLAN tunnels in single- and multi-homed deployments, simplifying troubleshooting and improving operational reliability.

For more details, refer to the [VxLAN Operation Administration Maintenance](#) section in the *OcNOS VxLAN Guide*, Release 7.0.0.

CLI-Script and CLI-Shell

Overview

The cli-script and cli-shell feature provides command automation and system command execution within the OcNOS command-line interface.

The cli-script function supports creation of script files that contain configuration mode and execution mode commands. These scripts can be executed in execution mode to apply the defined commands and store the resulting configuration on the system.

The cli-shell function enables execution of Linux bash commands directly from execution mode through the CLI.

Feature Characteristics

- Supports creation of cli-script files using the cli-script `file-name` command.
- Accepts configuration mode and execution mode commands as script input.
- Supports execution control through delay and message commands.
- Provides configurable behavior for error handling during script execution.
- Allows execution of linux bash commands using the exec-shell `linux command interface`.
- Stores cli-script files persistently on the file system.

Benefits

- Enables automation of operational and configuration workflows.
- Simplifies application of repetitive or grouped configuration changes.
- Reduces manual configuration effort and execution time.
- Provides controlled access to system-level commands from the CLI.

Limitations

- Editing an existing cli-script is not supported through the CLI.
- Modifying a script requires deleting and recreating the cli-script file.
- Built-in linux shell commands are not supported through the exec-shell interface.
- Improper use of cli-shell commands may affect system stability.

Configuration

CLI-Script Configuration

The cli-script feature allows the user to create a cli-script and add a set of commands to it, making it possible to apply a specific set of configurations at once when applying the cli-script. The main objective is to provide the creation of a cli-script in execution mode, with the cli-script `file-name` command, that enters the cli-script mode, and receives as input a series of commands. The name of the file has a limit of 128 characters and verifies invalid characters, such as >, <, *, among others.

1. Create a cli-script file and enter cli-script mode to define the sequence of commands.

```
OcNOS# cli-script TRANSLATION
```

This command creates a cli-script named `translation` and switches the CLI to cli-script mode.

2. Enter the configuration and execution mode commands that must be applied together when the script is executed.

Include all necessary commit commands within the script to ensure that configuration changes are applied.

```
OcNOS(cli-script)# configure terminal
OcNOS(cli-script)# netconf translation openconfig
OcNOS(cli-script)# commit
```



Note: The load-cli-script command does not perform an implicit commit. Any configuration commands included in the script must explicitly contain commit statements.

3. Exit cli-script mode and save the script contents to the system.

```
OcNOS(cli-script)# cli-script-end
```

This action saves the cli-script file and returns the CLI to execution mode.

4. Execute the saved cli-script to apply the defined commands.

```
OcNOS# load-cli-script TRANSLATION
OcNOS# show running-config netconf translation
```

The system executes each command in the script sequentially.

CLI-Shell Configuration

1. Execute a Linux bash command directly from execution mode.

```
OcNOS# exec-shell ip netns exec zebosfib0 ip addr show eth1
```

2. Execute an existing shell script from the system.

```
OcNOS# exec-shell /root/test_hello.sh
```

Configuration for Delay and Message Commands

1. Configure execution delay.

```
OcNOS# delay 5
OcNOS# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)#do delay 5
OcNOS(config)#
```

2. Display a custom message during execution.

```
OcNOS# message Test message
OcNOS# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
OcNOS(config)# do message Test message
```

Validation

Validate the creation of a cli-script with the following show commands available in CLI Execution mode:

1. Verify cli-script creation.

```
#show cli-script all
Filename      Last Modified
EVPN-MAINT-1  01-01-2010
TRANSLATION   02-01-2010
INT-EXEC-2    04-01-2020
```

2. Verify cli-script content.

```
OcNOS#show cli-script content all
cli-script test
configure terminal
interface xel
shutdown
commit
cli-script-end
cli-script test1
conf term
int xel
shutdown
commit
cli-script-end
```

Configuration Snapshot

```
OcNOS#show running-config extended
!
! Software version: EC_AS7315-30X-OcNOS-AGGR-NA-7.1.0.999- 01/20/2026 17:36:25
!
! Last configuration change at 16:04:58 UTC Thu Jan 22 2026 by root
!
!
netconf translation openconfig
!
service password-encryption
!
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!!
qos enable
!
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
!
ip vrf management
!
interface eth0
ip vrf forwarding management
ip address dhcp
!
exit
!
!
end
!
cli-script EVPN-MAINT-1
conf t
interface eth1
shutdown
commit
cli-script-end
cli-script TRANSLATION
conf
netconf translation openconfig
commit
cli-script-end
!
cli-script INT-EXEC-2
mtu 999
```

```
interface xe2
ip address 4.4.4.4/24
config terminal
interface xe1
shutdown
delay 20
no shutdown
commit
cli-script-end
#show cli-script EVPN-MAINT-1
conf t
interface eth1
shutdown
commit
OcNOS#show cli-script content all
cli-script EVPN-MAINT-1
conf t
interface eth1
shutdown
commit
cli-script-end
```

Implementation Examples

- CLI-Script allows the user to execute a sequence of commands, facilitating the execution of multiple tests automatically, saving each test and pasting it to other devices with the copy and paste options.
- Delay and message commands can be used inside cli-scripts to facilitate the understanding of the test execution process and what is being applied to the system in each moment.
- EXEC-Shell commands can help productivity by allowing the user to execute `shell` commands directly in CLI, without having the need to exit execution mode to access root and execute the command.

CLI-Script and CLI-Shell Commands

cli-script

Use this command to create a cli-script. The file contains a set of commands that can be applied together through an execution command.

When the command is executed, the file is created with the name specified by the user, and `cli-script` mode is accessed. In this mode, it is possible to add a list of commands.

Command Syntax

```
cli-script LINE
```

Parameters

LINE

Name of the cli-script to be created

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#cli-script test
(cli-script)#
OcNOS#no cli-script test
CLI-Script test deleted
```

cli-script line command

Use this command to add lines to a CLI script. This command does not require a prepend string, and CLI script mode accepts every string except the CLI script end string. Each line that is added is saved individually to the CLI script.

Command Syntax

LINE

Parameters

LINE

Commands to add to cli-script

Default

None

Command Mode

cli-script mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#cli-script test
OcNOS(cli-script)#configure terminal
OcNOS(cli-script)#interface xe2
OcNOS(cli-script)#shutdown
OcNOS(cli-script)#commit
```

cli-script-end

Use this command to exit `cli-script` mode and return to `execution` mode.

Command Syntax

```
cli-script-end
```

Parameters

None

Default

None

Command Mode

cli-script mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#cli-script test
(cli-script)#configure terminal
(cli-script)#interface xe2
(cli-script)#shutdown
(cli-script)#commit
(cli-script)#cli-script-end
OcNOS#
```


show cli-script

Use this command to display cli-script contents.

Command Syntax

```
show cli-script ( WORD | all )
```

Parameters

WORD

CLI-Script name

all

CLI-Script list

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#show cli-script all
Filename                Last Modified
test                    Wed Aug  6 01:41:20 2025
test1                   Wed Aug  6 01:43:30 2025
test2                   Wed Aug  6 01:42:08 2025
OcNOS#show cli-script test
interface xe2
shutdown
commit
```

load-cli-script

Use this command to apply a cli-script to the system. The application starts from Execution mode, accepting both configure and execution commands.

Command Syntax

```
load-cli-script LINE ( continue-on-error | stop-on-error | )
```

Parameters

LINE

CLI-Script name to be loaded from local

continue-on-error

Continue to process configuration and exec commands on error (default)

stop-on-error

Stop processing configuration and exec commands on error

Default

continue-on-error is the default.

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0.

Example

```
OcNOS#load-cli-script test stop-on-error
OcNOS#load-cli-script test continue-on-error
OcNOS#load-cli-script test
```

exec-shell

Use this command to execute Linux commands directly in CLI from Execution mode.

Blocked Shell commands

To ensure secure operation, exec-shell restricts execution of the following commands:

- bash
- dash
- gdb
- nano
- passwd
- sh
- vim
- vi
- yangcli

Command Syntax

```
exec-shell LINE
```

Parameters

LINE

Command to be executed.

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0.

Examples

```
OcNOS#exec-shell ip netns exec zebosfib0 ip addr show eth0
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc pfifo_fast state UP group default qlen 1000
link/ether 52:54:00:e0:68:44 brd ff:ff:ff:ff:ff:ff
altname enpls0
inet 192.168.100.3/24 brd 192.168.100.255 scope global eth0
valid_lft forever preferred_lft forever
inet6 fe80::5054:ff:fee0:6844/64 scope link
valid_lft forever preferred_lft forever
```

delay

Use this command to stop system execution for a specified duration in seconds. It functions like the `sleep` command in Linux and can be applied with the `do` option inside configuration mode.

Command Syntax

```
delay <0-1800>
```

Parameters

<0-1800>

Delay time interval (in seconds)

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#delay 10
<terminal waits for 10 seconds>
OcNOS(config)#do delay 10
<terminal waits for 10 seconds>
OcNOS(config)#
```

message

Use this command to display the message entered by the user in the terminal.

Command Syntax

```
message LINE
```

Parameters

LINE

Message to be displayed

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS##message Hello world!
Hello world!
OcNOS(config)#do message Hello world!
Hello world!
OcNOS(config)#
```

show running-config extended

Use this command to show both the running configuration and all CLI script content.

Command Syntax

```
show running-config extended
```

Parameters

None

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#show running-config extended?  
extended          Show cli-script extended  
extended-community-list  Extended-community-list  
<cr>  
OcNOS#show running-config extended
```

show cli-script content all

Use show cli-script content all to display the contents of all CLI scripts.

Command Syntax

```
show cli-script content all
```

Parameters

None

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS# show cli-script content all
cli-script test1
mtu 444
end
cli-script-end
cli-script test2
ip address 5.5.5.5/24
commit
cli-script-end
OcNOS#
```

no cli-script

Use this command to remove all the cli-script files present in the device.

Command Syntax

```
no cli-script (LINE | all)
```

Parameters

LINE

CLI-script name to be loaded from local

all

Delete all CLI-scripts

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#no cli-script test
CLI-Script removed: test
OcNOS#show cli-script all
Filename                               Last Modified
test3                                  Fri Jun 27 15:40:56 2025
test1                                  Fri Jun 27 15:30:01 2025
test2                                  Fri Jun 27 15:30:16 2025
OcNOS#no cli-script all
OcNOS#show cli-script all
Filename                               Last Modified
OcNOS#
```


copy running-config-ext <remote-location>

Use this command to export the current running configuration, including all active CLI configuration and script content, to a specified remote location such as a TFTP, FTP, or SCP server. This command enables backup of the system's current operational state for archival, troubleshooting, or migration purposes.

Command Syntax

```
copy running-config-ext (tftp TFTP-URL|ftp FTP-URL|scp SCP-URL|sftp SFTP-URL|http HTTP-URL) (vrf
(NAME|management) |)
```

Parameters**tftp TFTP-URL**

"upload files via tftp", "Enter URL tftp://server[:port]][/path/filename]"

ftp FTP-URL

"upload files via ftp", "Enter URL ftp://server[/path/filename]"

scp SCP-URL

"upload files via scp", "Enter URL scp://server[/path/filename]"

sftp SFTP-URL

"upload files via sftp", "Enter URL sftp://server[/path/filename]"

http HTTP-URL

"upload files via http", "Enter URL http://server[/path/filename]"

vrf NAME

"Specify VRF by name for the transfer"

vrf management

"Use management VRF for the transfer"

Default

None

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 7.0.0

Example

```
OcNOS#copy running-config-ext scp scp://root:root123@10.16.99.116/home/backup.txt vrf management
% Total      % Received % Xferd  Average Speed   Time    Time     Time  Current
Dload  Upload  Total    Spent    Left  Speed
100  2084    0      0  100   2084      0  11593  --:--:--  --:--:--  --:--:--  11642
100  2084    0      0  100   2084      0  11580  --:--:--  --:--:--  --:--:--  11580
Copy Success
OcNOS#
```

System Limits and Counters

Overview

The System Limits and Counters (Show and NetConf) feature enhances OcNOS operational visibility by providing direct access to system capacity and utilization data across key subsystems. It acts as a diagnostic and planning tool that consolidates hardware and software resource information into a single, consistent framework.

OcNOS monitors various resource categories internally, such as interfaces, VLANs, routing tables, MPLS labels, MAC tables, and protocol sessions. This data is accessible through both CLI and management interfaces, enabling operators and automation systems to understand resource consumption and remaining capacity.

The [show hardware-routing limits \(page 283\)](#) and [show hardware-l2vpn-instances limits \(page 285\)](#) commands, along with their corresponding NetConf or gNMI extensions, provide a structured, real-time view of hardware and service-level resource availability across all operational domains. These features simplify capacity validation, improve troubleshooting accuracy, and support readiness assessments before service deployment or scaling.

OcNOS aggregates data from multiple subsystems, involving system management, platform drivers, routing protocols, and hardware abstraction layers into a unified schema. The information is normalized so monitoring tools and network administrators can consistently interpret system capacity, regardless of hardware platform or ASIC implementation.

Beyond providing visibility, this also enhances operational predictability by enabling proactive checks on system headroom. It verifies whether sufficient resources are available for additional routes, MPLS labels, VLANs, or BGP sessions before configuration changes are made.

Using YANG-based models and JSON or JSON-IETF encoded responses, OcNOS smoothly integrates with external management systems, inventory platforms, and automation frameworks, ensuring consistent resource tracking across large-scale deployments.

Feature Characteristics

Data Organization

Information is grouped into the following functional categories:

- **System Limits:** Displays overall system-level capacities for routing/LPM, MPLS, VXLAN/VNIs, VLANs, VRFs, and TCAM resources. Provides a global resource usage summary across the device.
- **Layer 3 Counters:** Reports IPv4 and IPv6 route usage and maximum supported route entries. Reflects route scale across connected, static, and dynamic routing tables.
- **MPLS Counters:** Added gNMI support to retrieve MPLS counters as labels info, VPLS, and VPWS circuits.
- **EVPN Counters:** Added gNMI support to retrieve the number of EVPN MAC routes.
- **Protocol Counters:** Summarizes session counters for BGP, QoS, IS-IS, SLA, LAG, and ACL. Displays relevant information to these protocols.

Benefits

Unified Capacity View: Consolidates per-subsystem capacity information into a single display.

Operational Validation: Enables verification of resource availability before deploying services.

Proactive Monitoring: Helps anticipate resource exhaustion by tracking usage trends.

Automation-Ready: Data accessible via NetConf and gNMI for integration with Network Management System (NMS), Business Support System (BSS) or Operations Support System (OSS), or telemetry systems.

Platform-Agnostic Design: Abstracts hardware-specific details into normalized, comparable counters.

System Limits and Counters Limitation

Data Characteristics

- Read-only operational data; no configuration impact.
- Data is refreshed dynamically and time-stamped.
- Displays an instantaneous system snapshot, not a historical trend.
- SNMP access is not supported.

Security and Access

- CLI access is restricted to privileged operational modes.
- NetConf or gNMI retrieval allowed for authenticated management sessions only.

Dependencies

- Requires base system management and hardware driver integration.
- YANG data models must be enabled for NetConf queries.

Platform-specific

- Counter scope varies by platform and ASIC capability.
- Hardware system limits are available only on Qumran2 (Q2) series platforms using Broadcom DNX.
- Supports QoS queuing buffer depth per interface for protocol counters.

Encoding


- gNMI proto encoding is not supported for ACL and BGP xpaths.
- gNMI proto encoding is not supported for protocol counters with paths that have complex keys.

System Limits and Counters Configuration

This use case verifies the total number of IPv4 routes installed in the system and validates that the counter is consistent across CLI and NetConf or YANG interfaces. The objective is to confirm that the backend routing subsystem and YANG data model remain synchronized when routes are added or withdrawn dynamically.

Topology

Figure 7. Sample Topology



Use Case: Verify Total Number of IPv4 Routes Installed

1. Check the software and hardware information using the `show version` and `show system-information` commands. Confirms software version and hardware.
2. Apply the below configuration to routers and verify the setup.

```
#show running-config
!
!
service password-encryption
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
!
tfo Disable
errdisable cause stp-bpdu-guard
feature dns relay
ip dns relay
ipv6 dns relay
!
ip vrf management
!
ip vrf vrf1
!
interface eth0
 ip vrf forwarding management
 ip address dhcp
!
interface lo
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe1
!
interface xe2
 ip address 20.0.0.1/30
!
interface xe3
!
interface xe4
```

```

ip vrf forwarding vrf1
ip address 200.0.0.1/30
isis circuit-type level-1
ip router isis 1
!
interface xe5
!
interface xe6
!
exit
!
router ospf 1
network 20.0.0.0/30 area 0.0.0.0
!
router isis 1 vrf1
net 49.0001.0000.0000.0001.00
!
!
end

```

3. Verify the total number of installed IPv4 routes. The output shows: Total IPv4 routes (all VRFs): 11.

```

OcNOS#show ip route vrf all summary

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
IP routing table maximum-paths    : 8
Total number of IPv4 routes       : 6
Total number of IPv4 paths        : 6
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
ospf            4
Total          6
FIB             6

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

-----
IP routing table name is management(1)
-----
IP routing table maximum-paths    : 8
Total number of IPv4 routes       : 2
Total number of IPv4 paths        : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total          2
FIB             2

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

-----
IP routing table name is vrf1(2)
-----
IP routing table maximum-paths    : 8
Total number of IPv4 routes       : 3
Total number of IPv4 paths        : 3
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
isis            1
Total          3

```

```
FIB                3

ECMP statistics (active in ASIC):
  Total number of IPv4 ECMP routes : 0
  Total number of IPv4 ECMP paths  : 0

Total number of IPv4 routes (All VRFs) : 11
```

4. Remove routes (e.g., shutdown interface) and verify. The route count decreases (e.g., 6 routes total).

```
(config)#interface xe2
(config-if)#shutdown

#show ip route vrf all summary

.....
.....
ECMP statistics (active in ASIC):
  Total number of IPv4 ECMP routes : 0
  Total number of IPv4 ECMP paths  : 0

Total number of IPv4 routes (All VRFs) : 6
```

5. Re-enable the interface and recheck. The route count has increased back to 11.

```
(config)#no interface xe2
(config-if)#shutdown

#show ip route vrf all summary

.....
.....
-----
IP routing table name is vrf1(2)
-----
IP routing table maximum-paths : 8
Total number of IPv4 routes   : 3
Total number of IPv4 paths    : 3
Pending routes (due to route max reached): 0
Route Source   Networks
connected      2
isis           1
Total          3
FIB            3

ECMP statistics (active in ASIC):
  Total number of IPv4 ECMP routes : 0
  Total number of IPv4 ECMP paths  : 0

Total number of IPv4 routes (All VRFs) : 11
```

6. Add Static Routes and Verify Counter Increase: The multiple secondary loopback addresses and redistributed static routes increase the number of routes significantly. (example: 67 routes).

```
!
! Last configuration change at 13:34:17 UTC Tue Jun 03 2025 by root
!
!
service password-encryption
!
logging console disable
logging monitor disable
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
hardware-profile statistics ingress-acl enable
!
qos enable
```

```
!  
hostname R-A-7014  
tfo Disable  
errdisable cause stp-bpdu-guard  
feature dns relay  
ip dns relay  
ipv6 dns relay  
!  
ip vrf management  
!  
interface eth0  
  ip vrf forwarding management  
  ip address dhcp  
!  
interface ge0  
!  
interface ge1  
  ip address 30.0.0.1/30  
!  
interface ge2  
!  
interface ge3  
!  
interface ge4  
!  
interface ge5  
!  
interface ge6  
!  
interface ge7  
!  
interface ge8  
!  
interface ge9  
!  
interface ge10  
  ip ospf priority 10  
!  
interface ge11  
!  
interface lo  
  ip address 127.0.0.1/8  
  ip address 10.0.0.1/30 secondary  
  ip address 10.0.0.5/30 secondary  
  ip address 10.0.0.9/30 secondary  
  ip address 10.0.0.13/30 secondary  
  ip address 10.0.0.17/30 secondary  
  ip address 10.0.0.21/30 secondary  
  ip address 10.0.0.25/30 secondary  
  ip address 10.0.0.29/30 secondary  
  ip address 10.0.0.33/30 secondary  
  ip address 10.0.0.37/30 secondary  
  ip address 10.0.0.41/30 secondary  
  ip address 10.0.0.45/30 secondary  
  ip address 10.0.0.49/30 secondary  
  ip address 10.0.0.53/30 secondary  
  ip address 10.0.0.57/30 secondary  
  ip address 10.0.0.61/30 secondary  
  ip address 10.0.0.65/30 secondary  
  ip address 10.0.0.69/30 secondary  
  ip address 10.0.0.73/30 secondary  
  ip address 10.0.0.77/30 secondary  
  ip address 10.0.0.81/30 secondary  
  ip address 10.0.0.85/30 secondary  
  ip address 10.0.0.89/30 secondary  
  ip address 10.0.0.93/30 secondary  
  ip address 10.0.0.97/30 secondary  
  ip address 10.0.0.101/30 secondary
```

```
ip address 10.0.0.105/30 secondary
ip address 10.0.0.109/30 secondary
ip address 10.0.0.113/30 secondary
ip address 10.0.0.117/30 secondary
ip address 10.0.0.121/30 secondary
ip address 10.0.0.125/30 secondary
ip address 10.0.0.129/30 secondary
ip address 10.0.0.133/30 secondary
ip address 10.0.0.137/30 secondary
ip address 10.0.0.141/30 secondary
ip address 10.0.0.145/30 secondary
ip address 10.0.0.149/30 secondary
ip address 10.0.0.153/30 secondary
ip address 10.0.0.157/30 secondary
ip address 10.0.0.161/30 secondary
ip address 10.0.0.165/30 secondary
ip address 10.0.0.169/30 secondary
ip address 10.0.0.173/30 secondary
ip address 10.0.0.177/30 secondary
ip address 10.0.0.181/30 secondary
ip address 10.0.0.185/30 secondary
ip address 10.0.0.189/30 secondary
ip address 10.0.0.193/30 secondary
ip address 10.0.0.197/30 secondary
ip address 10.0.0.201/30 secondary
ip address 10.0.0.205/30 secondary
ip address 10.0.0.209/30 secondary
ip address 10.0.0.213/30 secondary
ip address 10.0.0.217/30 secondary
ip address 10.0.0.221/30 secondary
ip address 10.0.0.225/30 secondary
ip address 10.0.0.229/30 secondary
ip address 10.0.0.233/30 secondary
ip address 10.0.0.237/30 secondary
ip address 10.0.0.241/30 secondary
ip address 10.0.0.245/30 secondary
ip address 10.0.0.249/30 secondary
ip address 10.0.0.253/30 secondary
ipv6 address ::1/128
!
interface lo.management
 ip vrf forwarding management
 ip address 127.0.0.1/8
 ipv6 address ::1/128
!
interface xe12
!
interface xe13
!
interface xe14
 ip address 44.44.0.1/24
!
interface xe15
!
interface xe16
!
interface xe17
!
 exit
!
router bgp 100
 neighbor 20.0.0.2 remote-as 100
 neighbor 30.0.0.2 remote-as 100
!
 address-family ipv4 unicast
 redistribute static
 neighbor 20.0.0.2 activate
 neighbor 30.0.0.2 activate
```



```

exit-address-family
!
exit
!
line vty 0 16
exec-timeout 0 0
!
!

#show ip route summary

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
IP routing table maximum-paths    : 8
Total number of IPv4 routes       : 67
Total number of IPv4 paths        : 67
Pending routes (due to route max reached): 0
Route Source      Networks
connected         67
Total             67
FIB               67

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

```

7. Cross-verify YANG counters. YANG output shows the same route count (example: 69 — minor delta possible due to management VRFs or transient state).

```

#sget /ipi-rib:routing/global/counters

yangcli ocnos@127.0.0.1> sget /ipi-rib:routing/global/counters/total-routes-ipv4-vrf

RPC Data Reply 1 for session 30:

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <routing xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-rib">
      <global>
        <counters>
          <total-routes-ipv4-vrf>69</total-routes-ipv4-vrf>
        </counters>
      </global>
    </routing>
  </data>
</rpc-reply>

yangcli ocnos@127.0.0.1> sget /ipi-rib:routing/global/counters

Filling container /routing/global/counters:
RPC Data Reply 2 for session 30:

<rpc-reply xmlns="urn:ietf:params:xml:ns:netconf:base:1.0">
  <data>
    <routing xmlns="http://www.ipinfusion.com/yang/ocnos/ipi-rib">
      <global>
        <counters>
          <total-routes-ipv4-vrf>69</total-routes-ipv4-vrf>
          <total-routes-ipv6-vrf>4</total-routes-ipv6-vrf>
        </counters>
      </global>
    </routing>
  </data>
</rpc-reply>

```

```
(config-if)#do show ip route summ
% Warning: Executing the CLI from higher config mode level

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
IP routing table maximum-paths      : 8
Total number of IPv4 routes         : 67
Total number of IPv4 paths          : 67
Pending routes (due to route max reached): 0
Route Source      Networks
connected         67
Total             67
FIB               67

ECMP statistics (active in ASIC):
  Total number of IPv4 ECMP routes : 0
  Total number of IPv4 ECMP paths  : 0
```

The system accurately reports the total IPv4 route count through both CLI and NetConf or YANG interfaces, confirming backend and data model synchronization.

System Limits and Counters Implementation Example

Scenario 1: Resource Audit During Large-Scale Migration

During migration or network expansion, operators can periodically capture usage metrics to ensure resource growth aligns with expectations.

OcNOS consolidates routing table utilization and capacity, providing a clear snapshot of routing scalability during migration, with real-time reflection of control-plane resource allocation.

Scenario 2: Monitoring via Network Management Systems

External management tools can periodically poll the YANG path to maintain system-wide inventory and detect nearing resource limits.

OcNOS exposes these metrics through its NetConf or gNMI interfaces, enabling seamless integration with telemetry collectors and proactive alerting systems.

Scenario 3: Capacity Pre-Check Before Service Deployment

When provisioning a new service, such as an MPLS, EVPN, VXLAN, or Routing, operations teams can execute the [show hardware-capacity limits](#) and [show hardware-l2vpn-instances limits \(page 285\)](#) commands to ensure that available resources are sufficient for new configuration changes, reducing provisioning failures and service-impacting errors.

OcNOS provides a unified, hardware-aware query of all relevant resource pools, ensuring deployment teams can validate platform headroom without manual calculations or cross-referencing subsystem outputs.

System Limits and Counters Commands

show hardware-routing limits

The [show hardware-routing limits \(page 283\)](#) command displays the maximum supported system capacities and current utilization across operational domains.

Other show commands display partial limit and usage details per subsystem, such as `show interface summary`, `show vlan`, and `show bgp summary`. This command consolidates these separate outputs, providing a comprehensive summary across multiple domains.

show hardware-l2vpn-instances limits

The command [show hardware-l2vpn-instances limits \(page 285\)](#) displays the utilization of Virtual Network Identifier (VNI), Virtual Switching Instances (VSI), and Multicast Groups resources, categorized by the service type as E-LAN or E-LINE.

System Limits and Counters Revised Commands

The following commands are enhanced to include summary fields that display system-level totals. These additions improve visibility into configured resources across the system.

show ip vrf

The [show ip vrf \(page 287\)](#) command output includes the **Total Number of all VRFs** field, which displays the count of all configured VRFs in the system.

show access-lists summary

The [show access-list summary](#) command output includes the **Total ACEs configured on system** field, which displays the cumulative number of Access Control Entries (ACEs) across all access lists.

show vlan brief

The command output includes a **"Total Vlans"** field, which shows the total number of VLANs configured on the system. For more details, refer to the `show vlan brief` command section in the *OcNOS Layer 2 Guide*.

show ip route vrf all summary

The [show ip route \(page 288\)](#) command output includes the **Total number of IPv4 routes (All VRFs)** field, which displays the total number of IPv4 routes per VRF in the system.

show ipv6 route vrf all summary

The [show ipv6 route \(page 292\)](#) command output includes the **Total number of IPv6 routes (All VRFs)** field, which displays the total number of IPv6 routes per VRF in the system.

show interface <lag-if-name>

The [show interface \(page 294\)](#) command output includes the **Aggregator UP-Time** field, which displays the total UP duration for the aggregated interface.

show interface brief

The [show interface \(page 294\)](#) command output includes the **UP-Time** field, which displays the time duration for which the interface has remained in the UP state.

show hardware-routing limits

Use this command to display the maximum supported system capacities and current utilization across operational domains.

**Notes:**

- Only available on the DNX platforms.
- The values displayed in this command output estimate the available entries in the hardware. Because hardware tables are shared among multiple resources, the limit values may change based on which resources are currently in use.

Command Syntax

```
show hardware-routing limits
```

Parameters

None

Default

None

Command Mode

Privileged execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following is a sample output from this command.

```
OcNOS#show hardware-routing limits
```

```
Max Entries: Estimated max capacity. Depends on shared memory with other
              features, address range, prefix-length.
```

```
Used:        Number of entries retrieved from hardware.
```

```
Free Entries: Free entries calculated from max and used entries.
```

Database	Free Entries	%	Used Entries	Max Entries
LPM	2117619	0	13	2117632
IPv4 LPM	2117627	0	5	2117632
IPv6 LPM	1058812	0	4	1058816
MPLS labels POP	60289	0	3	60292
MPLS labels SWAP	715978	0	0	715978

Here is an explanation of the show command output fields.

Table 2. show hardware-routing limits output fields

Field	Description
Database	Hardware tables
Free Entries	Number of entries available in the hardware table.

Table 2. show hardware-routing limits output fields (continued)

Field	Description
Used Entries	Number of entries used by the resource.
Max Entries	Maximum number of entries supported in the hardware table.

show hardware-l2vpn-instances limits

Use this command to display the utilization of Virtual Network Identifier (VNI), Virtual Switching Instances (VSI), and Multicast groups resources, categorized by the service type as E-LAN or E-LINE.



Notes:

- This command is supported on Dune-based platforms.
- The displayed values are estimates of the hardware table entries available for each resource type.

Command Syntax

```
show hardware-l2vpn-instances limits
```

Parameters

None

Default

None

Command Mode

Privileged execution mode

Applicability

Introduced in OcNOS version 7.0.0.

Example

The following is a sample output from this command.

```
OcNOS#show hardware-l2vpn-instances limits
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               | Instances |                               | VSI |                               | M |
| Multicast Groups |          |                               |     |                               |   |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               |          | Used |          |          | Used |          |          |          |
| Service         | Used    |      |         |         |      |         |         |         |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                               | Free    |      | Max    | Free    |      | Max    | Free    |      |
|-----+-----+-----+-----+-----+-----+-----+-----+-----+
| % | Entries | Entries | % | Entries | Entries | Entries | % | Entries | Entries | Entries |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
E-LINE
```

```

VNI      14205      0      0      14205      28410      0      0      28410      90112      0      0
E-LAN
VNI      14205      0      0      14205      28410      0      0      28410      90112      0      0
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| Total      | 14205      0      0      14205      | 28410      0      0      28410      |
90112      0      0      90112      |
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
* 131 VSI entries are reserved for L3 interfaces. The real hardware capacity is 28541 VSI entries

```

Here is an explanation of the show command output fields.

Table 3. show hardware-l2vpn-instances limits output fields

Field	Description
Service	Identifies the service type, such as E-LINE VNI or E-LAN VNI.
Instances	Displays usage information for the service instances.
VSI	Shows utilization of Virtual Switching Instances (VSI) associated with the service.
Multicast Groups	Displays usage details for multicast group resources.
Free Entries	Number of hardware table entries currently available.
Used Entries	Number of hardware table entries currently in use.
Max Entries	Maximum number of entries supported in the hardware table.

show access-lists

Use this command to display access lists.

Command Syntax

```
show access-lists (NAME|) (expanded|summary|)
```

Parameters

NAME

Access-list name.

expanded

Expanded access-list.

summary

Summary of access-list.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3. Added the “Total ACEs configured on system” field in the show access-list summary command show output in OcNOS version 7.0.0.

Example

```
#show access-lists expanded
IP access list Iprule1
11 permit ip 30.0.0.1 0.0.0.255 172.124.0.2 0.0.0.255
default deny-all
MAC access list Macrule1
10 permit host 0000.1234.1234 any
default deny-all
IPv6 access list ipv6-acl-01
10 deny ahp 3ffe::/64 4ffe::/64
default deny-all

#show access-lists summary
IPV4 ACL Iprule1
statistics enabled
Total ACEs Configured: 1
Configured on interfaces:
xe3/1 - egress (Router ACL)
Active on interfaces:
xe1/3 - ingress (Router ACL)
MAC ACL Macrule1
statistics enabled
Total ACEs Configured: 0
Configured on interfaces:
Active on interfaces:
IPV6 ACL ipv6-acl-01
statistics enabled
Total ACEs Configured: 2
Configured on interfaces:
xe7/1 - ingress (Router ACL)
Active on interfaces:
Total ACEs configured on system: 3
```

show ip vrf

Use this command to display the routing information about VRFs.

Command Syntax

```
show ip vrf
show ip vrf WORD
```

Parameters

WORD

Virtual Routing and Forwarding name.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3. Added the “Total Number of all VRF's” field in the show output in OcNOS version 7.0.0.

Example

```
OcNOS#show ip vrf
VRF management, VRF ID: 1, FIB ID 1, MTU 1500
MPLS DSCP Preserve Disbaled (global)
Router ID: 10.16.179.120 (automatic)
Interfaces:
  eth0
  lo.management
!
Total Number of configured IP VRF's: 1
Total Number of all VRF's: 2
Maximum Number of VRF's: 4096
```

Name	Default RD
management	not set

show ip route

Use this command to display the IP routing table for a protocol or from a particular table.

When multiple entries are available for the same prefix, NSM uses an internal route selection mechanism based on protocol administrative distance and metric values to choose the best route. All best routes are entered into the FIB and can be viewed using this command. To display all routes (selected and not selected), use the **show ip route database** command.

Use this command to see all subnets of a specified network if they are present in the routing table. Use this command with mask information.

Command Syntax

```
show ip route A.B.C.D
show ip route (database|)
show ip route (database|) (bgp|connected|database|isis|fast-
reroute|interface|isis|kernel|mbgp|mstatic|next-hop|ospf|rip|static)
show ip route summary
show ip route vrf all summary
show ip route vrf WORD (database|)
show ip route vrf WORD (database|) (bgp|connected|isis|kernel|ospf|rip|static|summary)

show ip route mpls (summary|)
```

Parameters

A.B.C.D

Network in the IP routing table.

A.B.C.D/M

IP prefix <network>/<length>, for example, 35.0.0.0/8.

bgp

Border Gateway Protocol

connected

Connected (directly attached) routes.

database

Routing table database.

fast-reroute

Fast reroute repair paths.

interface

Routes learned or tied to a specific interface.

isis

IS-IS routing-protocol routes.

kernel

Kernel (local OS) routes.

mbgp

Multiprotocol BGP (e.g., VPN or EVPN) routes.

mstatic

Multicast static routes.

next-hop

Routes based on a specific next-hop address.

ospf

Open Shortest Path First routing-protocol routes.

rip

Routing Information Protocol routing-protocol routes.

static

Static routes

summary

Summarize all routes

vrf WORD

Routes for a specific Virtual Routing and Forwarding (VRF) instance named WORD.

vrf all

Routes for all VRF instances.

mpls

Shows the routing table entries resolved through MPLS next-hops (NHLFE).

mpls summary

Shows the summary of MPLS-resolved routes in the routing table.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3.

Added `mpls summary` parameter and Total number of IPv4 routes (All VRFs) field in the `show ip route vrf all summary` command in the OcNOS version 7.0.0.

Example

Displays all routes in the IP routing table database.

The recursive `via MPLS FTN-ID <ID>`, `nhlfe-ix <INDEX>` field specifies the MPLS next-hop used for recursive resolution.

show ip route database

```
OcNOS#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C   *> 1.1.1.1/32 is directly connected, lo, installed 00:04:51, last update 00:04:51 ago
O   1.1.1.1/32 [110/1] is directly connected, lo, installed 00:04:51, last update 00:04:51 ago
C   *> 1.2.0.0/24 is directly connected, xe5, installed 00:04:51, last update 00:04:51 ago
O   1.2.0.0/24 [110/1] is directly connected, xe5, installed 00:04:51, last update 00:04:51 ago
O   *> 2.2.2.2/32 [110/2] via 1.2.0.2, xe5, installed 00:04:24, last update 00:04:24 ago
B   *> 9.9.9.9/32 [20/0] via 2.2.2.2 (recursive via 1.2.0.2), installed 00:02:51, last update 00:02:51 ago
    *> [20/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1)
B   *> 10.10.10.10/32 [20/0] via 2.2.2.2 (recursive via 1.2.0.2), installed 00:02:51, last update 00:02:51 ago
    *> [20/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1)
S   *> 11.11.11.11/32 [1/0] via 2.2.2.2 (recursive via 1.2.0.2), installed 00:04:24, last update 00:04:51 ago
    *> [1/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1)
S   *> 12.12.12.12/32 [1/0] via 2.2.2.2 (recursive via 1.2.0.2), installed 00:04:24, last update 00:04:51 ago
    *> [1/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1)
S   *> 13.13.13.13/32 [1/0] via 2.2.2.2 (recursive via 1.2.0.2), installed 00:04:24, last update 00:04:51 ago
    *> [1/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1)
C   *> 127.0.0.0/8 is directly connected, lo, installed 00:05:45, last update 00:20:53 ago

Total number of IPv4 routes 16

Gateway of last resort is not set
```

show ip route mpls

```
OcNOS#show ip route mpls
IP Route Table for VRF "default"
B   9.9.9.9/32 [20/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1), installed 00:02:22, last update 00:02:22 ago
B   10.10.10.10/32 [20/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1), installed 00:02:22, last update 00:02:22 ago
S   11.11.11.11/32 [1/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1), installed 00:03:55, last update 00:04:22 ago
S   12.12.12.12/32 [1/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1), installed 00:03:55, last update 00:04:22 ago
S   13.13.13.13/32 [1/0] via 2.2.2.2 (recursive via MPLS FTN-ID 1, nhlfe-ix 1), installed 00:03:55, last update 00:04:22 ago

Gateway of last resort is not set
```

The Total number of routes `nhlfe` field shows all recursive routes resolved via MPLS next-hops. The `FIB` count field confirms how many MPLS-resolved routes are installed in hardware forwarding tables.

show ip route mpls summary

```
OcNOS#show ip route mpls summary

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
MPLS maximum-paths          : 1
Total number of IPv4 MPLS routes : 5
Total number of routes nhlfe   : 5
Pending routes (due to route max reached): 0
Route Source      Networks
static            3
bgp                2
Total             5
FIB               5
```

Here are the explanations for the show command output fields.

Table 4. show ip route mpls summary

Field	Description
MPLS maximum-paths	Maximum number of ECMP paths supported for MPLS recursive routes.
Total number of IPv4 MPLS routes	Total IPv4 routes using MPLS recursive resolution.
Total number of routes nhlfe	Total routes resolved through MPLS next-hops (NHLFE).
Pending routes	Routes awaiting installation due to resource or configuration limits.
Route Source	Distribution of routes by protocol (static, BGP, etc.).
FIB	Number of MPLS-resolved routes installed in hardware forwarding tables.

```
OcNOS#show ip route vrf all summary

-----
IP routing table name is Default-IP-Routing-Table(0)
-----
IP routing table maximum-paths : 8
Total number of IPv4 routes    : 1
Total number of IPv4 paths     : 1
Pending routes (due to route max reached): 0
Route Source      Networks
connected         1
Total             1
FIB               1

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

-----
IP routing table name is management(1)
-----
IP routing table maximum-paths : 8
Total number of IPv4 routes    : 2
Total number of IPv4 paths     : 2
Pending routes (due to route max reached): 0
Route Source      Networks
connected         2
Total             2
FIB               2

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
```

```

Total number of IPv4 ECMP paths   : 0

-----
IP routing table name is red(2)
-----
IP routing table maximum-paths   : 8
Total number of IPv4 routes      : 2
Total number of IPv4 paths       : 2
Pending routes (due to route max reached): 0
Route Source      Networks
connected         2
Total             2
FIB               2

ECMP statistics (active in ASIC):
Total number of IPv4 ECMP routes : 0
Total number of IPv4 ECMP paths  : 0

VRF FIB Route Limits:
Configured Route Limit   : 1000
Utilization Percentage   : 0 %
Action upon reaching limit: stop-install
Warning Threshold        : 80 %
Exceeds Threshold        : No

Total number of IPv4 routes (All VRFs) : 5

```

show ipv6 route

Use this command to display the IP routing table for a protocol or from a particular table, including database entries known by NSM. When multiple entries are available for the same prefix, NSM uses an internal route selection mechanism based on protocol administrative distance and metric values to choose the best route. The best routes in the FIB can be viewed using **show ipv6 route**.

Command Syntax

```

show ipv6 route vrf WORD (database|)
show ipv6 route vrf WORD (database|) (bgp|connected|isis|kernel|ospf|rip|static|summary)
show ipv6 route (database)
show ipv6 route (database) (bgp|connected|isis|kernel|ospf|rip|static)
show ipv6 route X:X::X:X
show ipv6 route X:X::X:X/M
show ipv6 route summary
show ipv6 route vrf all summary

```

Parameters

X:X::X:X

Network in the IP routing table.

X:X::X:X/M

Prefix <network>/<length>, e.g., 35.0.0.0/8

all

All IPv6 routes

bgp

Border Gateway Protocol.

connected

Connected.

database

IPv6 routing table database.

isis

IS-IS.

IFNAME

Interface name

kernel

Kernel.

ospf

Open Shortest Path First.

rip

Routing Information Protocol.

static

Static routes.

summary

Summarize all routes

vrf WORD

Routes from a Virtual Routing and Forwarding instance.

vrf all

Routes for all VRF instances.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3. Added the **Total number of IPv6 routes (All VRFs)** field to the `show ipv6 route vrf all summary` display output in OcNOS version 7.0.0.

Examples

See [route codes and modifiers](#) and [route entry output details](#) tables for an explanation of the codes and fields in the output.

```
#show ipv6 route
Codes: K - kernel route, C - connected, S - static, R - RIPng, O - OSPFv3,
       I - IS-IS, B - BGP, > - selected route, * - FIB route, p - stale info.
C> * ::1/128 is directly connected, lo
C> * 3ffe:1::/48 is directly connected, eth1
C> * 3ffe:2:2::/48 is directly connected, eth2
```

```
OcNOS#show ipv6 route vrf all summary

-----
IPv6 routing table name is Default-IPv6-Routing-Table(0)
-----
IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 1
Total number of IPv6 paths       : 1
Pending routes (due to route max reached): 0
```

```

Route Source    Networks
connected       1
Total           1
FIB              1

ECMP statistics (active in ASIC):
  Total number of IPv6 ECMP routes   : 0
  Total number of IPv6 ECMP paths    : 0

-----
IPv6 routing table name is management(1)
-----

IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB              2

ECMP statistics (active in ASIC):
  Total number of IPv6 ECMP routes   : 0
  Total number of IPv6 ECMP paths    : 0

-----
IPv6 routing table name is red(2)
-----

IPv6 routing table maximum-paths : 8
Total number of IPv6 routes      : 2
Total number of IPv6 paths       : 2
Pending routes (due to route max reached): 0
Route Source    Networks
connected       2
Total           2
FIB              2

ECMP statistics (active in ASIC):
  Total number of IPv6 ECMP routes   : 0
  Total number of IPv6 ECMP paths    : 0

Total number of IPv6 routes (All VRFs) : 5

```

show interface

Use this command to display interface configuration and status information.

Command Syntax

```

show interface (IFNAME|)
show interface brief (IFNAME|)

```

Parameters

IFNAME

Interface name

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3. Added the UP-Time field to the `show interface <LAG-IFNAME>` and `show interface brief` commands output in OcNOS version 7.0.0.

Example

```
#show interface xe1/1
Interface xe1/1
  Scope: both
  Flexport: Breakout Control Port (Active): Break Out Enabled
  Hardware is ETH Current HW addr: ecf4.bb6e.934b
  Physical:ecf4.bb6e.934b Logical:(not set)
  Port Mode is access
  Interface index: 5001
  Metric 1 mtu 1500 duplex-full(auto) link-speed 1g(auto)
  PHY Link Training: Disabled
  PHY Dfe: Enabled
  PHY Unreliable LOS: Disabled
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
DHCP client is disabled.
Last Flapped: 2016 Nov 05 22:40:23 (00:19:25 ago)
Statistics last cleared: 2016 Nov 05 04:49:55 (18:09:53 ago)
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 256 bits/sec, 0 packets/sec
RX
  unicast packets 39215813 multicast packets 0 broadcast packets 0
  input packets 39215813 bytes 2666662432
  jumbo packets 0
  runs 0 giants 0 CRC 0 fragments 0 jabbers 0
  input error 0
  input with dribble 0 input discard 0
  Rx pause 0
TX
  unicast packets 38902 multicast packets 437 broadcast packets 0
  output packets 437 bytes 28018
  jumbo packets 0
  output errors 0 collision 0 deferred 0 late collision 0
  output discard 0
  Tx pause 0

OcNOS#show interface brief
.....
-----
Port-channel Type PVID Mode Status Reason Speed  UP-Time
Interface
-----
-----
po10          AGG      1 trunk   up    none   100g   00:00:38
sa10          AGG      1 trunk   down  PD      0      00:00:00
.....
```

show interface <LAG-IFNAME>

```
OcNOS#show interface po10
Interface po10
Hardware is AGG Current HW addr: 5c07.5851.cd03
Physical:(Not Applicable) Logical:(not set)
Aggregator UP-Time: 00:00:38
```

```

Port Mode is trunk
.....

OcNOS#show interface sa10
Interface sa10
Hardware is AGG Current HW addr: 5c07.5851.cd04
Physical:(Not Applicable) Logical:(not set)
Aggregator UP-Time: 00:00:00
Port Mode is trunk
.....

```

Here is the explanation of the show command output fields.

Table 5. show interface output details

Field	Description
Scope	Interface can be used for communication within the device and outside the device (Both).
Flexport	Specifies whether the ports has Breakout capabilities or is a Non-Control Port.
Breakout Control Port (Active)	Specifies whether Breakout is active or disabled.
Hardware is ETH Current HW addr	The MAC address of the interface.
Physical	Displays the physical MAC address of the interface.
Logical	Displays the logical MAC address (if any) of the interface.
Aggregator UP-Time	Shows the total UP duration for the aggregated interface.
Port Mode	Displays the port mode: Router, VLAN access, switch, or trunk.
Interface index	Index number, Metric, MTU size, duplex-full (auto) or half-duplex, minimum link speed in gigabits, and if the interface is up, broadcasting, and multicasting.
PHY Link Training	Displays the status of physical link training,
PHY Dfe	Displays the status of physical digital feedback equalizer.
PHY Unreliable LOS	Displays the status of physical unreliable loss of signal.
VRF Binding	Show whether the interface is VRF bound and (if bound) with what VRF, if Label Switching is enabled or disabled, and if a virtual circuit is configured.
DHCP client	The state of the DHCP client – whether this interface is connected to a DHCP server.
Last Flapped	Date and time when the interface last flapped.
Statistics last cleared	Date and time when the interface's statistics were cleared.
5 minute input rate	Input rate in bits/second and packets/second
5 minute output rate	Output rate in bits/second and packets/second
RX	Counters for unicast packets, multicast packets, broadcast packets, input packets, bytes, jumbo packets, runts, giants, CRC errors, fragments, jabbers, input errors, input with dribble input discards, and receive pause.
TX	Counters for unicast packets, multicast packets, broadcast packets, output packets, bytes, jumbo packets, output errors, collisions, differed packets, input late collisions, output discards, and transmit pause.


```
#show interface brief xe51
```

```
Codes: ETH - Ethernet, LB - Loopback, AGG - Aggregate, MLAG - MLAG Aggregate
FR - Frame Relay, TUN -Tunnel, PBB - PBB Logical Port, VP - Virtual Port
CVP - Channelised Virtual Port, METH - Management Ethernet, UNK- Unknown
ED - ErrDisabled, PD - Protocol Down, AD - Admin Down, IA - InActive
PD(Min L/B) - Protocol Down Min-Links/Bandwidth
OTD - Object Tracking Down
DV - DDM Violation, NA - Not Applicable
NOM - No operational members, PVID - Port Vlan-id
Ctl - Control Port (Br-Breakout/Bu-Bundle)
```

```
-----
Ethernet  Type      PVID  Mode      Status Reason  Speed Port Ch #  Ctl Br/Bu  Loopbk
Interface
-----
xe51      ETH        --    routed    down   OTD     10g   --          No          No
```

System Limit Counters Troubleshooting

Show Output Displays Blank or Partial Results

Possible Cause

- System drivers or the hardware abstraction layer did not initialize counter data.
- The underlying subsystem has not yet reported usage statistics.

Action

- Wait for the system to complete initialization after reboot.
- Confirm that relevant features (e.g., MPLS, EVPN, or L3 routing) are enabled and active.
- Re-run the command after a few seconds to verify the updated output.

gNMI Returns “Unsupported Encoding” or Missing Fields

Possible Cause

- Proto encoding is not supported for certain xpaths (ACL, BGP, or complex key-based data models).
- The client requested a non-supported encoding type.

Action

- Use JSON or JSON-IETF encoding for the affected paths.
- Validate the YANG path in the `Get` request matches the supported schema.
- Reissue the request with supported encoding formats.

API Retrieval Fails for Specific Resource Paths

Possible Cause

- Requested data path not supported by the platform or missing YANG model capability.
- The YANG model for the feature is not loaded or not enabled in the management subsystem.

Action

- Verify that the YANG model is present in the package directory.
- Ensure NetConf or gNMI service is enabled.
- Confirm the correct namespace and hierarchy are used in the `Get` query.

System Limit Counters Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Term or Acronym	Description
Access Control Entry (ACE)	An individual rule within an access control list (ACL) used to permit or deny specific traffic.
Access Control List (ACL)	A set of access rules applied to interfaces or packets for filtering network traffic.
Application Programming Interface (API)	A set of protocols and tools that allow external systems or applications to interact with OcNOS using programmatic methods such as NetConf or gNMI.
Border Gateway Protocol (BGP)	A routing protocol used to exchange routing and reachability information among autonomous systems on the Internet.
Broadcom DNX Series (DNX)	A family of high-capacity network switching ASICs that support scalable networking features and counters.
Ethernet Virtual Private Network (EVPN)	A BGP-based control plane solution that provides Layer 2 and Layer 3 VPN services over an MPLS or VXLAN core.
gRPC Network Management Interface (gNMI)	A protocol used to access and manage network configuration and operational data using YANG models over gRPC transport.
JSON or JSON-IETF	Data encoding formats used in management APIs for structured data exchange. JSON-IETF aligns with IETF YANG model representation standards.
L3 Counters	Counters that display Layer 3 resource utilization, including IPv4 and IPv6 route entries.
Multiprotocol Label Switching (MPLS)	A packet-forwarding technology that uses labels to make forwarding decisions.
Network Configuration Protocol (NetConf)	A protocol used to install, manipulate, and delete configurations of network devices using YANG-based data models.
Quality of Service (QoS)	Mechanisms that manage bandwidth allocation, delay, and packet prioritization in network traffic.
Ternary Content Addressable Memory (TCAM)	A high-speed memory used for packet classification, ACLs, and forwarding table lookups.
Virtual Routing and Forwarding (VRF)	A logical routing instance that allows multiple routing tables to coexist on the same device.
Yet Another Next Generation (YANG)	A data modeling language used to model configuration and state data for network management protocols such as NetConf and gNMI.

Key Term or Acronym	Description
Business Support System (BSS)	Software applications that support service providers business operations such as billing, product management, and customer management.
NetConf	A REST-like API mechanism in OcNOS that retrieves operational data or system state using YANG-based NetConf or gNMI Get operations.
Network Management System (NMS)	A centralized system that monitors, manages, and controls network devices and services.
Operations Support System (OSS)	Tools and systems used by service providers to manage network operations, provisioning, and fault management.
System Limits	Hardware and software resource capacities that define the maximum supported instances of configurable objects (e.g., interfaces, VLANs, MAC entries).
Telemetry	A framework that continuously exports operational data from network devices to external collectors or monitoring systems.

| LAYER 2 OR LAYER 3 OVERLAY NETWORKING

Enhancements in overlay networking expand OcNOS support for EVPN, MPLS, and VXLAN-based deployments. These updates improve scalability, simplify multi-tenant connectivity, and optimize traffic engineering across Layer 2 and Layer 3 overlays.

Flow-Aware Transport for EVPN over MPLS	301
Overview	301
FAT Prerequisites	301
FAT Configuration	305
FAT Implementation Example	311
FAT Revised Command	312
FAT Glossary	314
Recursive Next-hop Resolution with MPLS Next-hop	316
Overview	316
Limitation	316
Configuration	316
Implementation Example	329
Revised Commands	329
Troubleshooting	330
Glossary	332

Flow-Aware Transport for EVPN over MPLS

Overview

Flow-Aware Transport (FAT) Labels introduce entropy into the MPLS label stack to improve Equal-Cost Multipath (ECMP) forwarding in EVPN networks. By inserting a FAT label as the bottom-of-stack (BOS) below the EVPN service label, per-flow hashing is enabled across core routers, ensuring efficient traffic distribution without requiring transport-layer changes. This label is independent of any modifications in the Transport Label, ensuring end-to-end entropy delivery.



Note: BGP-based FAT label signaling, as defined in RFC 7432, is not supported.

Feature Characteristics

- FAT Label can be configured on Provider Edge (PE) routers for Unicast EVPN traffic, but validation occurs only on the ingress or egress interfaces of Transit (P) nodes. It does not apply to BUM traffic.
- Supports both single-homed and multihomed deployments.
- Applies to EVPN E-LAN, E-LINE, and E-Tree services.
- Maintains compatibility with existing Control Word (CW) support.

Benefits

- **Improved Load Balancing:** Ensures even distribution of traffic across ECMP paths.
- **Higher Throughput:** Maximizes available bandwidth utilization by avoiding congestion hotspots.
- **Lower Latency:** Reduces queuing and congestion delays in the core.
- **Data Center and Large-Scale EVPN Ready:** Addresses east-west traffic patterns in leaf-spine and redundant mesh topologies.
- **No Transport Dependency:** Entropy is provided by FAT label insertion, not by core transport capabilities.

FAT Prerequisites

Before configuring FAT in [Figure 8](#) topology, ensure the following:

- MPLS, LDP, RSVP, and IS-IS (or OSPF) are operational across all PEs and core routers.
- Loopback interfaces for each router are reachable over the IGP.
- Control-word and flow-label options are supported together for unicast traffic.
- EVPN MPLS and Multihoming are globally enabled.
- BGP EVPN sessions are established among all PEs.
- Port-channels or physical links toward the core have MPLS, LDP, RSVP, and IGP enabled.
- [Core Routers \(P1 and P2\)](#) form the MPLS transport fabric between PE nodes.
- [Access Switches \(SW-1 and SW-2\)](#) are configured with VLAN trunks toward PEs.

PE Nodes

[PE1 and PE2 Nodes](#) and [PE3 and PE4 Nodes](#) provide Active-Standby EVPN connectivity and originate the flow-label insertion for FAT-based ECMP or LAG load balancing.

Enable EVPN MPLS and Multihoming

Enables EVPN MPLS operation with active-standby redundancy between PEs.

```
!  
hardware-profile filter evpn-mpls-mh enable  
!  
evpn mpls enable  
evpn mpls multihoming enable  
!  
evpn mpls vtep-ip-global 1.1.1.1  
!  
evpn etree enable scenario-2  
!
```

Configure Core-Facing Interfaces

Core-facing interfaces carry MPLS, IS-IS, LDP, and RSVP signaling for underlay transport.

```
!  
interface po10  
  description connected-to-p1  
  load-interval 30  
  ip address 12.1.1.1/24  
  mtu 9216  
  label-switching  
  mpls ldp-igp sync ospf holddown-timer 10  
  ip router isis 100  
  enable-ldp ipv4  
  mpls ldp-igp sync-delay 5  
  enable-rsvp  
!  
interface po20  
  description connected-to-p2  
  load-interval 30  
  ip address 13.1.1.1/24  
  mtu 9216  
  label-switching  
  mpls ldp-igp sync ospf holddown-timer 10  
  ip router isis 100  
  enable-ldp ipv4  
  mpls ldp-igp sync-delay 5  
  enable-rsvp  
!
```

Configure Route Distinguishers and Route Targets

Defines route-target import or export for EVPN services.

```
!  
mac vrf elan_vrf116  
  rd 1.1.1.1:116  
  route-target both 116:116  
!  
  
!  
mac vrf eline_vrf216
```

```
rd 1.1.1.1:216
route-target both 216:216
!

!
mac vrf etree_vrf316
rd 1.1.1.1:316
route-target both 316:316
!
```

Physical and Port-Channel Interfaces

Prepares physical interfaces and link aggregation for connectivity and redundancy. Configure similar link aggregation in PEs for consistent topology.

- **channel-group:** Aggregates multiple physical interfaces into one logical link for redundancy and higher throughput.
- **mode active:** Configures LACP active mode.
- **speed (g):** Ensures high-speed core links.

```
!
interface ce5
channel-group 1000 mode active
!
interface xe10
channel-group 10 mode active
!
interface xe11
channel-group 10 mode active
!
interface xe12
speed 10g
channel-group 20 mode active
!
interface xe26
channel-group 20 mode active
!
```

Configure LDP and RSVP

Enables MPLS label distribution and RSVP synchronization with entropy label capability for FAT forwarding.

```
!
router ldp
targeted-peer ipv4 2.2.2.2
exit-targeted-peer-mode
targeted-peer ipv4 3.3.3.3
exit-targeted-peer-mode
targeted-peer ipv4 4.4.4.4
exit-targeted-peer-mode
transport-address ipv4 1.1.1.1
neighbor 2.2.2.2 auth md5 password plain-text test1
!
router rsvp
!
rsvp-trunk PE1-PE3 ipv4
to 3.3.3.3
!
rsvp-trunk PE1-PE4 ipv4
to 4.4.4.4
!
rsvp-trunk PE1-PE2 ipv4
```

```
to 2.2.2.2
!
```

Configure OSPF for MPLS Reachability

OSPF advertises IP routes so MPLS LSPs can be established.

```
!
router ospf 100
  ospf router-id 1.1.1.1
  fast-reroute keep-all-paths
  shutdown
  bfd all-interfaces
  network 1.1.1.1/32 area 0.0.0.0
  network 12.1.1.0/24 area 0.0.0.0
  network 13.1.1.0/24 area 0.0.0.0
!
```

IGP and MPLS Setup

Provides routing and MPLS TE capabilities with FAT support. Configure identical IS-IS or MPLS setup on PEs to enable full TE and FAT ECMP forwarding.

- **mpls traffic-eng router-id:** Identifies router for RSVP and TE.
- **capability cspf:** Enables CSPF for TE path computation.
- **bfd all-interfaces:** Fast failure detection for all interfaces.

```
!
router isis 100
  is-type level-1
  ignore-lsp-errors
  lsp-gen-interval 5
  spf-interval-exp level-1 50 2000
  metric-style wide
  mpls traffic-eng router-id 1.1.1.1
  mpls traffic-eng level-1
  capability cspf
  dynamic-hostname
  bfd all-interfaces
  net 49.0001.0100.0000.1111.00
!
```

Configure BGP for MPLS Services

Supports L3VPN or service connectivity over MPLS or FAT.

```
!
router bgp 65010
  neighbor PG1 peer-group
  neighbor PG1 remote-as 65010
  neighbor 2.2.2.2 peer-group PG1
  neighbor 3.3.3.3 peer-group PG1
  neighbor 4.4.4.4 peer-group PG1
!
address-family l2vpn evpn
  neighbor PG1 activate
exit-address-family
!
```

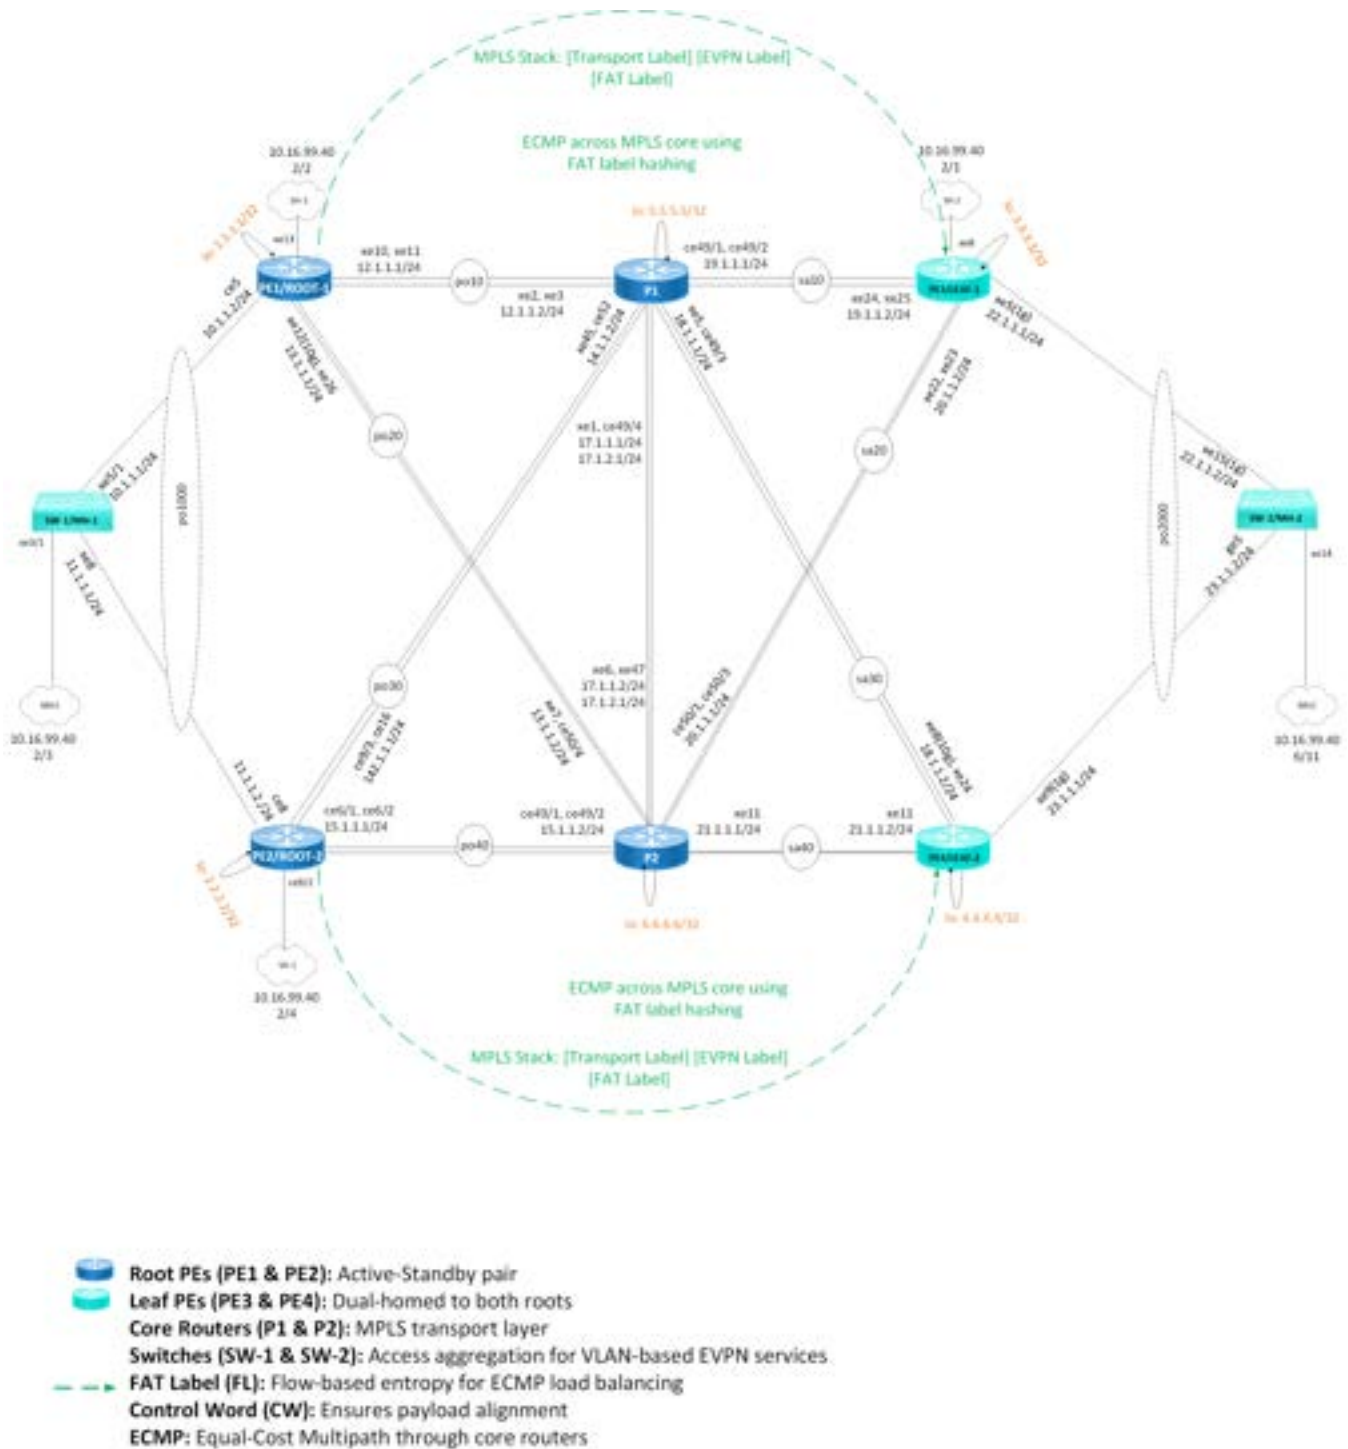
FAT Configuration

This section describes a sample use case topology illustrating the FAT label deployment in an EVPN MPLS network. It explains configuring the flow-label on Provider Edge (PE) routers, its operation across the MPLS core, and validation steps to ensure proper functionality and traffic distribution.

Topology

The topology represents an EVPN Active-Standby network where PE1 and PE2 nodes are connect to the PE3 and PE4 nodes through intermediate core routers (P1 and P2) over an MPLS transport network. Each PE is connected to customer edge (CE) switches (SW-1 and SW-2) that provide end-user or service access. Each PE is configured with the `flow-label both` option, which inserts a FAT label beneath the EVPN service label (Bottom of Stack). This FAT label introduces per-flow entropy, allowing Equal-Cost Multipath (ECMP) load balancing across the MPLS core. As a result, traffic between CE switches and the PEs (PE1 and PE2) is efficiently distributed across multiple core paths. The configuration is compatible with control-word support and applies to both EVPN E-LINE and E-LAN services, enhancing bandwidth utilization, reducing congestion, and improving overall network performance.

Figure 8. Flow-Label Enabled EVPN Active-Standby Topology



Note: Before configuration meet all [FAT Prerequisites](#) (page 301).

PE Nodes Configuration



Note: Configure the flow-label option on all participating PE nodes for each EVPN service (E-LINE or E-LAN). This ensures consistent FAT label insertion for a given EVPN instance, enabling end-to-end per-flow entropy and ECMP load balancing across the MPLS core.

EVPN E-LAN and E-LINE Services

Defines EVPN services with optional control-word and flow-label for per-flow ECMP and sequencing. EVPN IDs must match PE's configuration for multihoming.

- **flow-label:** Enables per-flow entropy to improve ECMP load sharing.
- **control-word:** Ensures pseudowire sequencing compatibility.
- **xconnect:** Defines E-LINE pseudowire connection to remote PE.

E-LAN Services

```
!  
evpn mpls id 116 control-word flow-label  
host-reachability-protocol evpn-bgp elan_vrf116  
!
```

E-LINE Services

```
!  
evpn mpls id 5016 xconnect target-mpls-id 6016 control-word flow-label  
host-reachability-protocol evpn-bgp eline_vrf216  
!
```

E-Tree Services

```
!  
evpn mpls id 316 control-word flow-label  
host-reachability-protocol evpn-bgp etree_vrf316  
!
```

Map VLANs to VPN Instances

Associates VLANs or sub-interfaces with EVPN services and multihoming.

- **evpn multi-homed system-mac:** Ensures consistent MAC address for multihomed VLANs.
- **map vpn-id:** Maps access VLAN to the respective EVPN service.

```
!  
interface po1000  
  evpn multi-homed system-mac 0000.1212.1313  
!  
  
!  
interface po1000.116 switchport  
  description evpn-elan-services-between-pe1-pe2-pe3-pe4-with-CW-and-FL  
  encapsulation dot1q 116  
  load-interval 30  
  access-if-evpn
```

```
    map vpn-id 116
!
!
interface po1000.216 switchport
description evpn-eline-services-MH-MH-with-CW-and-FL
encapsulation dot1q 216
load-interval 30
access-if-evpn
map vpn-id 5016
!
!
interface po1000.316 switchport
description evpn-etree-services-MH-MH-with-CW-and-FL
encapsulation dot1q 316
load-interval 30
access-if-evpn
map vpn-id 316
!
```

FAT Validation

After configuration deployment, the following checks and validations confirm that routing, MPLS, EVPN services, and FAT load-balancing are functioning correctly.

IGP (OSPF or IS-IS) Neighbor Verification

The show commands `show ip ospf neighbor` and `show clns neighbor` confirms that OSPF or IS-IS adjacency is established between PEs and routers. Verifies neighbor state is Full (for OSPF) or Up (for IS-IS), indicating proper routing protocol operation. Ensures loopbacks and links are advertised for MPLS and EVPN tunnels.

MPLS Control Plane Verification (LDP or RSVP)

The show commands `show ldp session` and `show rsvp session` confirms that LDP and RSVP sessions are operational. Ensures MPLS label distribution is complete for all PE-CE and PE-PE links. Verifies RSVP-TE tunnels are ready for traffic engineering and FAT ECMP.

EVPN Control Plane and Tunnel Verification

The following show commands:

- Confirms EVPN session establishment and MPLS tunnels between PEs.
- Checks EVPN service reachability, pseudowire integrity, and MAC address learning.
- Ensures flow-label and control-word settings are active, enabling FAT ECMP load balancing.
- Monitors tunnel counts and traffic statistics to validate correct EVPN distribution.

```
show evpn mpls tunnel summary
show evpn mpls tunnel
show evpn mpls tunnel label
show evpn mpls mac-table
show evpn mpls xconnect
show evpn mpls xconnect tunnel
show evpn mpls xconnect tunnel label
show evpn etree-leaf brief
show bgp l2vpn evpn
show evpn mpls route-count
```

```
show evpn mpls counters <network | evpn-id> <egress | ingress>
```

Post-AC Interface Flap Verification

Steps

- Flap the AC interface (shutdown or no shutdown).
- Verify traffic continuity end-to-end across EVPN services.
- Confirm load-balancing is occurring for services configured with flow-label.

Verification

- Ensures EVPN services maintain redundancy and failover works as expected.
- Confirms flow-label-based per-flow ECMP is distributing traffic evenly.

Packet-Level Verification

Steps

- Capture traffic across EVPN or VPWS tunnels. In this example Wireshark tool capture is used.
- Inspect the MPLS label stack in each packet:
 - Transport Label (Label 0) – Represents IPv4 Explicit Null; ensures the packet is carried with QoS preserved.
 - Service Label (EVPN Label) – Identifies the EVPN ELINE or ELAN instance associated with the payload.
 - Flow Label (FAT Label) – Provides per-flow entropy for ECMP load balancing.
 - Bottom-of-Stack (BoS) bit = 1 – Confirms the last label before payload (Ethernet frame or IP packet).
 - Control-word bit = set – Verifies correct pseudowire encapsulation.

Verification

- Confirms MPLS label stack integrity for each packet.
- Validates label order: Transport → Service → Flow (FAT).
- Ensures ECMP load balancing using the FAT label.
- Confirms the control word and BoS bit are correctly set.
- Verifies that the data plane encapsulation matches the configured EVPN services.

Figure 9. Sample Packet Capture

capture_on_g1_ce45_2_tx_withLpcap

File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help

Apply a display filter ... <Ctrl-/>

No.	Time	Source	Destination	Protocol	Length	Info
1	0.000000	00:00:00_03:a3:70	00:00:00_03:b3:70	0x88b5	104	Local Experimental Ethertype 1
2	0.000169	00:00:00_03:a3:c6	00:00:00_03:b3:c6	0x88b5	104	Local Experimental Ethertype 1
3	0.000577	00:00:00_03:a3:c8	00:00:00_03:b3:c8	0x88b5	104	Local Experimental Ethertype 1
4	0.000635	00:00:00_03:a3:c9	00:00:00_03:b3:c9	0x88b5	104	Local Experimental Ethertype 1
5	0.000798	00:00:00_03:a4:1d	00:00:00_03:b4:1d	0x88b5	104	Local Experimental Ethertype 1
6	0.017434	00:00:00_03:a4:6f	00:00:00_03:b4:6f	0x88b5	104	Local Experimental Ethertype 1
7	0.017542	00:00:00_03:a4:71	00:00:00_03:b4:71	0x88b5	104	Local Experimental Ethertype 1
8	0.017597	00:00:00_03:a4:72	00:00:00_03:b4:72	0x88b5	104	Local Experimental Ethertype 1
9	0.026060	00:00:00_03:a4:73	00:00:00_03:b4:73	0x88b5	104	Local Experimental Ethertype 1
10	0.026125	00:00:00_03:a4:7c	00:00:00_03:b4:7c	0x88b5	104	Local Experimental Ethertype 1
11	0.026181	00:00:00_03:a4:7d	00:00:00_03:b4:7d	0x88b5	104	Local Experimental Ethertype 1
12	0.034784	00:00:00_03:a4:ca	00:00:00_03:b4:ca	0x88b5	104	Local Experimental Ethertype 1
13	0.035012	00:00:00_03:a5:20	00:00:00_03:b5:20	0x88b5	104	Local Experimental Ethertype 1
14	0.043565	00:00:00_03:a5:22	00:00:00_03:b5:22	0x88b5	104	Local Experimental Ethertype 1
15	0.043623	00:00:00_03:a5:23	00:00:00_03:b5:23	0x88b5	104	Local Experimental Ethertype 1
16	0.043679	00:00:00_03:a5:2c	00:00:00_03:b5:2c	0x88b5	104	Local Experimental Ethertype 1
17	0.043734	00:00:00_03:a5:7c	00:00:00_03:b5:7c	0x88b5	104	Local Experimental Ethertype 1
18	0.043789	00:00:00_03:a5:7d	00:00:00_03:b5:7d	0x88b5	104	Local Experimental Ethertype 1
19	0.052263	00:00:00_03:a5:7f	00:00:00_03:b5:7f	0x88b5	104	Local Experimental Ethertype 1
20	0.052320	00:00:00_03:a5:84	00:00:00_03:b5:84	0x88b5	104	Local Experimental Ethertype 1

> Frame 1: 104 bytes on wire (832 bits), 104 bytes captured (832 bits)

> Ethernet II, Src: EdgecoreNetw_d1:76:79 (b8:6a:97:d1:76:79), Dst: EdgecoreNetw_9f:42:1f (d8:77:ce:9f:42:1f)

MultiProtocol Label Switching Header, Label: 0 (IPv4 Explicit-Null), Exp: 0, S: 0, TTL: 63

0000 0000 0000 0000 0000 = MPLS Label: IPv4 Explicit-Null (0)

..... 000. = MPLS Experimental Bits: 0

..... 0 = MPLS Bottom Of Label Stack: 0

..... 0011 1111 = MPLS TTL: 63

MultiProtocol Label Switching Header, Label: 24976, Exp: 0, S: 0, TTL: 64

0000 0110 0001 1001 0000 = MPLS Label: 24976 (0x06190)

..... 000. = MPLS Experimental Bits: 0

..... 0 = MPLS Bottom Of Label Stack: 0

..... 0100 0000 = MPLS TTL: 64

MultiProtocol Label Switching Header, Label: 576598, Exp: 0, S: 1, TTL: 0

1000 1100 1100 0101 0110 = MPLS Label: 576598 (0x8cc56)

..... 000. = MPLS Experimental Bits: 0

..... 1 = MPLS Bottom Of Label Stack: 1

..... 0000 0000 = MPLS TTL: 0

Ethernet II, Src: 00:00:00_03:a3:70 (00:00:00:03:a3:70), Dst: 00:00:00_03:b3:70 (00:00:00:03:b3:70)

> Destination: 00:00:00_03:b3:70 (00:00:00:03:b3:70)

> Source: 00:00:00_03:a3:70 (00:00:00:03:a3:70)

Type: 802.1Q Virtual LAN (0x8100)

[Stream index: 1]

802.1Q Virtual LAN, PRI: 0, DEI: 0, ID: 115

000. = Priority: Best effort (default) (0)

...0 = DEI: Ineligible

.... 0000 0111 0011 = ID: 115

Type: Local Experimental Ethertype 1 (0x88b5)

> Data (60 bytes)

Example EVPN Tunnel Status (PE1)

- **Status = Installed:** EVPN tunnel is active and ready.
- **FAT = set:** Flow-label is enabled for per-flow ECMP load balancing.
- Confirms all configured tunnels are active, indicating full connectivity across PEs.

E-LAN

```
PE1#show evpn mpls tunnel
EVPN-MPLS Network tunnel Entries
Source           Destination      Status           Up/Down          Update           evpn-id          Local-
Leaf  Remote-Leaf Ext-Color  FAT
```

```

=====
1.1.1.1      3.3.3.3      Installed    02:27:11    01:16:48    116      ---
   ---      ---      set
.
.
1.1.1.1      2.2.2.2      Installed    02:27:11    01:16:48    116      ---
   ---      ---      set
.
.
1.1.1.1      4.4.4.4      Installed    02:27:11    01:16:48    116      ---
   ---      ---      set

Total number of entries are 75

```

E-LINE

```

PE1#show evpn mpls xconnect tunnel
EVPN-MPLS Network tunnel Entries
Source      Destination      Status      Up/Down      Update      local-evpn-id remote-
evpn-id Ext-Color FAT
=====
1.1.1.1      2.2.2.2      Installed    03:28:13    01:16:49      5016      6016
   ---      set
1.1.1.1      3.3.3.3      Installed    03:28:13    01:16:49      5016      6016
   ---      set
1.1.1.1      4.4.4.4      Installed    03:28:13    01:16:49      5016      6016
   ---      set
.
.
.

Total number of entries are 12

```

E-Tree

```

PE1#show evpn etree-leaf brief

Leaf
Interface Ifindex  Vnid      Leaf
status  label
-----
po1000.100 0x1f400064 100      Root    ----
po1000.101 0x1f400065 101      Leaf     16
.
.
.
po1000.316 0x1f4000c8 200      Root    ----

Total number of entries are 22

```

FAT Implementation Example

In a data center leaf-spine fabric, multiple equal-cost paths exist between leaf switches (PEs) across the spine (core). Without FAT labels, EVPN traffic uses a single entropy-less service label, limiting load distribution. With FAT labels:

- Each flow receives a unique entropy label.
- Core P-routers hash on the FAT label, distributing traffic across all spine links.
- This achieves near-equal bandwidth utilization and avoids congestion.

EVPN Service Inactive

- Verify FAT label configuration consistency across PEs.
- Ensure control-word settings match between peers.

1. Traffic Not Balanced

- Confirm FAT label is present at BOS in MPLS stack using packet captures.
- Check core router hashing algorithms to ensure they use MPLS labels for entropy.

2. Interoperability Issues

- If the peer does not support FAT labels, disable static FAT configuration to maintain connectivity.

FAT Revised Command

The [evpn mpls id \(page 313\)](#) command includes the parameter `flow-label` to enable FAT label insertion under EVPN over MPLS. When configured, the router pushes a static FAT label below the EVPN service label, providing per-flow entropy for ECMP load balancing across the MPLS core.

evpn mpls id

Use this command to set the VPN identifier to create an EVPN MPLS tunnel.



Note: To set the VPN identifier for an E-LAN, use the `evpn mpls id` command. For creating an E-LINE/XConnect, use the `evpn mpls xconnect` command with source and target identifiers.

Use the `no` parameter of this command to delete the EVPN MPLS ID for the MPLS tunnel.

Command Syntax

```
evpn mpls id <1-16777215> (| xconnect target-mpls-id <1-16777215>) (|control-word) (|etree-leaf)
(|flow-label)
no evpn mpls id <1-16777215>
```

Parameters

evpn mpls id <1-16777215>

Specifies the EVPN-MPLS tenant identifier. This is a numeric value ranging from 1 to 16777215.

xconnect target-mpls-id <1-16777215>

Enables E-LINE Xconnect. Specifies the target EVID for E-LINE Xconnect. This is a numeric value within the range from 1 to 16777215.

control-word

Enables control-word egress or ingress options for the given EVPN Instance (E-LAN or E-LINE).

etree-leaf

(Optional) Configures the device as a leaf node within the E-Tree topology.

flow-label

Inserts a [Flow-Aware Transport \(FAT\) label](#) for per-flow hashing to improve ECMP traffic distribution.

Command Mode

Configure mode

Applicability

Introduced in OcNOS version 3.0.

Introduced `xconnect target-mpls-id <1-16777215>` parameter in the OcNOS version 4.0, `control-word` parameter in the OcNOS version 6.0.0, `etree-leaf` parameter in the OcNOS version 6.5.1, and `flow-label` parameter in the OcNOS version 7.0.0.

Example

```
(config)#evpn mpls id 10
(config-evpn-mpls)#exit

(config)#no evpn mpls id 10

(config)#evpn mpls id 100 xconnect target-mpls-id 200
(config-evpn-mpls)#exit

(config)#no evpn mpls id 100

(config)#evpn mpls id 300 xconnect target-mpls-id 400 control-word
```

```
(config-evpn-mpls)#exit
(config)#no evpn mpls id 300
```

Use the following command to configure the leaf node as an E-Tree leaf in a MPLS EVPN network.

```
(config)#evpn mpls id 10 etree-leaf
(config-evpn-mpls)#exit
```

The following configuration shows how to enable the `flow-label` parameter under different EVPN MPLS ID instances:

```
!
evpn mpls id 10 xconnect target-mpls-id 15 flow-label
  host-reachability-protocol evpn-bgp vrf2
!
evpn mpls id 20 xconnect target-mpls-id 25 control-word flow-label
  host-reachability-protocol evpn-bgp vrf2
!
evpn mpls id 100 flow-label
  host-reachability-protocol evpn-bgp vrf1
!
evpn mpls id 200 control-word flow-label
  host-reachability-protocol evpn-bgp vrf1
!
```

FAT Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Terms or Acronym	Description
Attachment Circuit (AC)	The physical or logical interface connecting a customer network to a Provider Edge (PE) router.
Bidirectional Forwarding Detection (BFD)	A protocol used to detect link failures rapidly between forwarding engines.
Bottom of Stack (BoS)	The last label in an MPLS label stack; when set to 1, it identifies the end of the label stack.
Control Word (CW)	A 4-byte field added before the payload to preserve packet sequencing and detect packet loss in pseudowires.
Equal-Cost Multipath (ECMP)	A routing technique that distributes traffic across multiple equal-cost paths to optimize bandwidth usage.
Ethernet LAN (E-LAN)	An EVPN service type that provides multipoint Layer 2 connectivity among multiple PEs.
Ethernet Line (E-LINE)	An EVPN service type that provides point-to-point Layer 2 connectivity between two PEs.
Ethernet VPN (EVPN)	A BGP-based control-plane mechanism for delivering Layer 2 and Layer 3 VPN services over MPLS or IP networks.
Flow-Aware Transport (FAT)	A mechanism using MPLS flow labels to achieve per-flow load balancing across ECMP paths.
Flow Label (FL)	A label in the MPLS stack that introduces per-flow entropy for balanced

Key Terms or Acronym	Description
	distribution across multiple paths.
Interior Gateway Protocol (IGP)	A routing protocol (such as OSPF or IS-IS) used within a service provider's network.
LLabel Distribution Protocol (LDP)	A protocol that distributes MPLS labels between routers to establish label-switched paths (LSPs).
Label Switched Path (LSP)	A unidirectional path established through an MPLS network, identified by labels at each hop.
MMultiprotocol Label Switching (MPLS)	A data-carrying technique that forwards packets based on short labels rather than long network addresses.
Provider Edge (PE) Router	A router located at the edge of the provider network that connects to customer sites.
Provider (P) Router	A core router in the provider network that forwards MPLS traffic between PEs.
Pseudowire (PW)	A virtual connection that emulates a point-to-point link over an MPLS core for Layer 2 services.
RResource Reservation Protocol (RSVP)	A signaling protocol used to reserve resources and establish label-switched paths in MPLS-TE networks.

Recursive Next-hop Resolution with MPLS Next-hop

Overview

Recursive Next-hop Resolution with MPLS Next-hop enables OcNOS to resolve recursive BGP and Static routes using MPLS transport paths instead of IP next-hops. This feature allows seamless packet forwarding through an MPLS core where BGP is not active, supporting scalable BGP-Free Core deployments. The system recognizes available MPLS Label Switched Paths (LSPs) as valid forwarding options, ensuring reachability across label-switched domains without requiring additional routing protocols in the backbone.

Feature Characteristics

- Supports route resolution over LDP, RSVP-TE, and Segment Routing (SR-MPLS) tunnels.
- Automatically updates route resolution when the underlying MPLS transport changes.
- Integrates with MPLS ECMP, FRR, and TI-LFA for resiliency and load balancing.
- Displays MPLS-resolved recursive routes in standard routing table views.
- Maintains consistent behavior for both BGP and Static recursive routes.
- Operates transparently without additional configuration changes to the routing process.

Benefits

Enables BGP-Free Core Operation: Backbone routers forward traffic using MPLS without running BGP.

Simplified Architecture: Reduces control-plane load and configuration complexity.

Improved Resiliency: Leverages MPLS protection mechanisms for faster failover.

Optimized Traffic Flow: Supports equal-cost load sharing over multiple MPLS paths.

Operational Visibility: Provides clear differentiation of MPLS-resolved routes for easier troubleshooting.

Limitation

- BGP multipath is not supported for BGP-Free Core deployments.
- If multiple BGP paths exist for the same prefix, the RIB selects only one path for MPLS next-hop resolution.
- Equal-Cost Multipath (ECMP) and Fast Reroute (FRR) are provided only by the underlying MPLS transport for that single resolved path.

Configuration

This section describes the Recursive Next-hop Resolution with MPLS Next-hop feature, illustrating its operation with a sample topology, configuration examples, and validation in an [MPLS LDP](#) and [SR-MPLS](#) deployments.

Configuration with MPLS LDP


Topology (MPLS LDP)

In this topology:

- PE1 or PE2 are Provider Edge routers connected to customer devices (CE1 or CE2).

- P1 or P2 are core Provider routers.
- LDP is enabled on all PE and P routers to exchange MPLS labels and establish LSPs.
- OSPF and IS-IS are used as the IGP to advertise reachability between LDP peers.
- BGP is used between PE routers for VPN route exchange (VRFs).

Figure 10. Recursive Next-hop Resolution with MPLS Next-hop in MPLS LDP Deployment



Enable MPLS LDP on Core and PE Routers

- LDP requires a router ID (typically loopback).
- LDP is enabled per interface connecting MPLS-capable peers.
- OSPF or IS-IS should be running to provide IGP reachability for LDP session establishment.

```
router ldp
  router-id 2.2.2.2
  !
  !
interface lo
  ip address 2.2.2.2/32 secondary
  !
interface xe14
  load-interval 30
  ip address 30.1.1.2/24
  label-switching
  ip ospf network point-to-point
  enable-ldp ipv4
  !
interface xe16.80
  encapsulation dot1q 80
  ip address 80.1.1.1/24
  ip ospf network point-to-point
  !
```

Configure IGP (OSPF or IS-IS)

- IGP ensures all LDP peers can discover each other.
- BFD is optional for fast detection of failures.

```
!
router ospf 1
  ospf router-id 2.2.2.2
  bfd all-interfaces
  network 2.2.2.2/32 area 0.0.0.0
  network 30.1.1.0/24 area 0.0.0.0
  !
```

Configure BGP (for VPNs or PE-to-PE routes)

- BGP exchanges prefixes between PE routers.
- MPLS labels are used to forward traffic between PE routers over the LDP-established LSPs.

```

!
router bgp 64000
  bgp router-id 2.2.2.2
  neighbor 8.8.8.8 remote-as 64000
  neighbor 8.8.8.8 update-source lo
!
address-family ipv4 unicast
  redistribute connected
  neighbor 8.8.8.8 activate
exit-address-family
!
!

```

Validation (MPLS LDP)**Verify LDP Sessions**

- **Expected output:** Session state = OPERATIONAL
- Confirms label exchange between LDP peers (e.g., PE1 ↔ P1).

```

PE1#show mpls ldp session
! [execution timestamp : 2025 Oct 09 16:25:19]
Codes: m - MD5 password is not set/unset.
       g - GR configuration not set/unset.
       t - TCP MSS not set/unset.
       Session has to be cleared manually

```

Code	Peer IP Address	IF Name	My Role	State	KeepAlive	UpTime	AutoCreate
	4.4.4.4	xe14	Passive	OPERATIONAL	30	00:40:43	No

Verify IP Routes

- Check that all networks (PE, P, CE) are reachable.
- Routes learned via OSPF or ISI-S or BGP should point to next-hop over MPLS.

```

PE1#show ip route
! [execution timestamp : 2025 Oct 09 16:25:08]
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       * - candidate default

```

IP Route Table for VRF "default"

```

C      2.2.2.2/32 is directly connected, lo, installed 1d14h08m, last update 1d14h08m ago
O      4.4.4.4/32 [110/2] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O      6.6.6.6/32 [110/3] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O      8.8.8.8/32 [110/4] via 30.1.1.4, xe14, installed 00:33:50, last update 00:33:50 ago
O      11.1.1.0/24 [110/2] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O      12.1.1.0/24 [110/3] via 30.1.1.4, xe14, installed 00:33:59, last update 00:33:59 ago
O      26.10.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:33:50, last update 00:33:50 ago
C      30.1.1.0/24 is directly connected, xe14, installed 00:40:53, last update 00:40:53 ago
C      80.1.1.0/24 is directly connected, xe16.80, installed 1d14h08m, last update 1d14h08m ago
B      81.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:41, last
update 00:33:41 ago
C      127.0.0.0/8 is directly connected, lo, installed 1d14h09m, last update 1d14h09m ago

```

```
O          222.1.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:33:50, last update 00:33:50 ago

Gateway of last resort is not set
```

Verify MPLS Labels (ILM or FTN)

- ILM table shows labels received from peers.
- FTN table shows outgoing labels and interfaces for forwarding.

```
PE1#show mpls ilm-table
! [execution timestamp : 2025 Oct 09 16:25:11]
Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
       K - CLI ILM, T - MPLS-TP, s - Stitched ILM
       S - SNMP, L - LDP, R - RSVP, C - CRLDP
       B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
       O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
       P - SR Policy,      U - unknown, UPStr - upstream

ILM-ECMP LDP: Enabled, SR: Enabled
Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF   Nexthop   pri   Algo-Num   UpTime   UPStr   peers
-----
B>  vrf100           2       26305     Nolabel     N/A       N/A       N/A
      Yes  N/A       00:40:49
B>  vrf200           1       26304     Nolabel     N/A       N/A       N/A
      Yes  N/A       00:40:49

PE1#show mpls forwarding-table
! [execution timestamp : 2025 Oct 09 16:25:13]
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Enabled, SR: Enabled
Code   FEC           FTN-ID   Nhlfe-ID   Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           Nexthop   Algo-Num   UpTime
-----
L>  4.4.4.4/32      1        7          -          -      -          -
      -           N/A      00:02:04
      Yes  0           xe14      Yes      30.1.1.4    -      -          -
L>  6.6.6.6/32      2        10         -          -      -          -
      -           N/A      00:02:04
      Yes  24960        xe14      No       30.1.1.4    -      -          -
L>  8.8.8.8/32      3        12         -          -      -          -
      -           N/A      00:02:04
      Yes  24961        xe14      No       30.1.1.4    -      -          -
L>  11.1.1.0/24     4        7          -          -      -          -
      -           N/A      00:40:37
      Yes  0           xe14      Yes      30.1.1.4    -      -          -
L>  12.1.1.0/24     5        15         -          -      -          -
      -           N/A      00:40:37
      Yes  24962        xe14      No       30.1.1.4    -      -          -
L>  26.10.1.0/24    6        17         -          -      -          -
      -           N/A      00:40:37
      Yes  24963        xe14      No       30.1.1.4    -      -          -
L>  222.1.1.0/24    10       36         -          -      -          -
      -           N/A      00:33:54
      Yes  24965        xe14      No       30.1.1.4    -      -          -
```

Example from PE1:

FEC	Out-Label	Out-Intf	Nexthop
6.6.6.6/32	24960	xe14	30.1.1.4
8.8.8.8/32	24961	xe14	30.1.1.4

Verify MPLS Recursive Route Resolution on PE1

The `show ip route database` and `show ip route mpls` outputs confirm that PE1 resolves BGP routes via MPLS Label Switched Paths (LSPs) established through LDP.

- The BGP route 81.1.1.0/24 (customer network behind PE2) is learned from the BGP peer 8.8.8.8 (PE2).
- The next hop 8.8.8.8 is not directly connected; it is recursively resolved via the IGP route to 30.1.1.4 (toward P1).
- The additional line (`recursive via MPLS FTN-ID 3, nhlfex 12`) indicates that PE1 uses an MPLS forwarding entry (FTN-ID 3) for this BGP route, proving that traffic destined for the remote CE network (81.1.1.0/24) is label-switched through the MPLS core rather than IP-routed.

```

PE1#show ip route database
! [execution timestamp : 2025 Oct 09 16:25:26]
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
C    *> 2.2.2.2/32 is directly connected, lo, installed 1d14h08m, last update 1d14h08m ago
O    2.2.2.2/32 [110/1] is directly connected, lo, installed 00:34:27, last update 00:34:27 ago
O    *> 4.4.4.4/32 [110/2] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
O    *> 6.6.6.6/32 [110/3] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
O    *> 8.8.8.8/32 [110/4] via 30.1.1.4, xe14, installed 00:34:08, last update 00:34:08 ago
B    8.8.8.8/32 [200/0] via 8.8.8.8 inactive, installed 00:33:59, last update 00:33:59 ago
O    *> 11.1.1.0/24 [110/2] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
O    *> 12.1.1.0/24 [110/3] via 30.1.1.4, xe14, installed 00:34:17, last update 00:34:17 ago
B    12.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago
O    *> 26.10.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:34:08, last update 00:34:08 ago
B    26.10.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago
C    *> 30.1.1.0/24 is directly connected, xe14, installed 00:41:11, last update 00:41:11 ago
O    30.1.1.0/24 [110/1] is directly connected, xe14, installed 00:34:27, last update 00:34:27
ago
C    *> 80.1.1.0/24 is directly connected, xe16.80, installed 1d14h08m, last update 1d14h08m ago
O    80.1.1.0/24 [110/1] is directly connected, xe16.80, installed 1d14h08m, last update 1d14h08m
ago
B    *> 81.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago
      *>                [200/0] via 8.8.8.8 (recursive via MPLS FTN-ID 3, nhlfex 12)
C    *> 127.0.0.0/8 is directly connected, lo, installed 1d14h09m, last update 1d14h09m ago
O    *> 222.1.1.0/24 [110/4] via 30.1.1.4, xe14, installed 00:34:08, last update 00:34:08 ago
B    222.1.1.0/24 [200/0] via 8.8.8.8 (recursive via 30.1.1.4), installed 00:33:59, last update
00:33:59 ago

Total number of IPv4 routes 20

Gateway of last resort is not set

```

```
PE1#show ip route mpls
```



```
! [execution timestamp : 2025 Oct 09 16:25:31]
IP Route Table for VRF "default"
B      81.1.1.0/24 [200/0] via 8.8.8.8 (recursive via MPLS FTN-ID 3, nhlfe-ix 12), installed
00:34:04, last update 00:34:04 ago

Gateway of last resort is not set
```

Test End-to-End Reachability (MPLS LDP)

- Ping or traceroute should reach CE2 via MPLS LSPs (PE1 → P1 → P2 → PE2).
- MPLS labels are swapped at core nodes (P1, P2) automatically.

```
PE1#ping 81.1.1.1
Press CTRL+C to exit
PING 81.1.1.1 (81.1.1.1) 100(128) bytes of data.
108 bytes from 81.1.1.1: icmp_seq=1 ttl=64 time=0.709 ms
108 bytes from 81.1.1.1: icmp_seq=2 ttl=64 time=0.583 ms
108 bytes from 81.1.1.1: icmp_seq=3 ttl=64 time=0.564 ms

--- 81.1.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2059ms
rtt min/avg/max/mdev = 0.564/0.618/0.709/0.064 ms
PE1#
PE1#
PE1#traceroute 81.1.1.1
traceroute to 81.1.1.1 (81.1.1.1), 30 hops max, 60 byte packets
 1  81.1.1.1 (81.1.1.1)  0.901 ms  0.852 ms  0.981 ms
```

Observed Behavior (MPLS LDP)


- **LDP Sessions:** PE1 establishes LDP with P1, P1 with P2, and P2 with PE2.
- **MPLS Forwarding:** Labels assigned to each destination allow PE1 to reach PE2's CE network (81.1.1.0/24) via core LSP.
- **BGP and MPLS Interaction:** BGP routes for customer networks are advertised and attached to labels. PE1 forwards traffic using MPLS labels without IP routing lookup at core routers.
- **Validation confirms:**
 - IP routes are correct (show ip route).
 - MPLS label tables are populated (show mpls ilm-table and show mpls forwarding-table).
 - End-to-end connectivity works (ping and traceroute).

Configuration with Segment Routing MPLS (SR-MPLS)

Topology (SR-MPLS)

In this topology:

- PE1 or PE2 are Provider Edge routers connected to customer devices (CE1 or CE2).
- P1 or P2 are core Provider routers.
- SR-MPLS is enabled on PE and P routers to provide MPLS transport using IGP-advertised Prefix-SIDs (no LDP sessions are required).
- OSPF is used as the IGP to advertise reachability and distribute Prefix-SID mappings across the network.
- BGP is enabled between PE routers to exchange customer prefixes; recursive next-hop resolution uses the SR-MPLS transport path for forwarding.

Figure 11. Recursive Next-hop Resolution with MPLS Next-hop in SR-MPLS Deployment

Configure Loopback and Prefix-SID

Each router must have a loopback router-id and a Prefix-SID used by SR-MPLS to build labeled transport paths.

- `prefix-sid index` assigns the SID index for the loopback prefix.
- `no-php` disables penultimate hop popping for the Prefix-SID label.

```

!
interface lo
  ip address 127.0.0.1/8
  ip address 1.0.0.1/32 secondary
  ipv6 address ::1/128
  prefix-sid index 101 no-php
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!

```

Enable Label Switching on Core-facing Interfaces

Enable MPLS label switching on interfaces that participate in the SR-MPLS transport. This ensures SR labels are pushed or swapped or popped through the MPLS transport path.

P1

```

!
interface ge6
  load-interval 30
  ip address 10.1.12.1/24
  mtu 9216
  label-switching
  ip router isis 100
!
interface ge7
!
interface ge8
  load-interval 30
  ip address 10.1.101.2/24
  mtu 9216
  label-switching
  ip router isis 100
!

```

Configure IGP with SR enabled (OSPF)

OSPF provides reachability and distributes Prefix-SID mappings required for SR-MPLS transport.

- `segment-routing mpls` command enables SR-MPLS in the OSPF process.
- (Optional) `fast-reroute per-prefix ti-lfa` command enables fast reroute using TI-LFA.

```
!  
router ospf 100  
  ospf router-id 1.0.0.1  
  network 1.0.0.1/32 area 0.0.0.0  
  network 10.1.12.0/24 area 0.0.0.0  
  network 10.1.101.0/24 area 0.0.0.0  
  segment-routing mpls  
  fast-reroute per-prefix ti-lfa area 0.0.0.0  
!
```

Configure BGP (PE-to-PE routes)

BGP exchanges customer prefixes between PE routers. With this feature, recursive BGP next-hops can be resolved via SR-MPLS transport.

```
!  
router bgp 100  
  bgp router-id 1.1.1.1  
  neighbor 2.2.2.2 remote-as 100  
  neighbor 2.2.2.2 update-source lo  
  neighbor 2.2.2.2 advertisement-interval 0  
  !  
  address-family ipv4 unicast  
  redistribute connected  
  neighbor 2.2.2.2 activate  
  exit-address-family  
  !  
exit  
!
```

Validation (SR-MPLS)

Verify Segment Routing State

- **Expected output:** SR State: SR_ENABLED and Operational state: enabled
- Confirms Segment Routing is enabled for the OSPF process and SR-MPLS transport is operational.

```
PE1#show ip ospf segment-routing state
```

```
OSPF process 100 Segment-Routing:  
SR State: SR_ENABLED  
SRGB Start: 16000,   SRGB Range: 8000  
SRLB Start: 14080,   SRLB Range: 1920  
Operational state: enabled
```

```
PE2#show ip ospf segment-routing state
```

```
OSPF process 100 Segment-Routing:  
SR State: SR_ENABLED  
SRGB Start: 16000,   SRGB Range: 8000  
SRLB Start: 14080,   SRLB Range: 1920  
Operational state: enabled
```

Verify Prefix-SID Mapping

- **Expected output:** Mapping table lists loopback prefixes with SID indexes for all nodes.
- Confirms Prefix-SIDs are advertised and learned across the SR domain.

```
PE1#show ip ospf segment-routing mapping-table active  
OSPF process ID: 100
```

Conflict Resolution Policy: Quarantine

Prefix	SID Index	Range	Flags
1.0.0.1/32	101	1	
1.0.0.2/32	102	1	
1.1.1.1/32	111	1	
2.2.2.2/32	222	1	

Number of mapping entries in Active Table: 4

```
PE2#show ip ospf segment-routing mapping-table active
OSPF process ID: 100
Conflict Resolution Policy: Quarantine
```

Prefix	SID Index	Range	Flags
1.0.0.1/32	101	1	
1.0.0.2/32	102	1	
1.1.1.1/32	111	1	
2.2.2.2/32	222	1	

Number of mapping entries in Active Table: 4

Verify IP Routes

Check that all networks (PE, P, CE) are reachable. Routes learned via OSPF or BGP must provide reachability to the remote PE loopback and remote customer subnet.

```
PE1#show ip route
```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
 O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
 ia - IS-IS inter area, E - EVPN,
 v - vrf leaked
 * - candidate default

IP Route Table for VRF "default"

```
O      1.0.0.1/32 [110/2] via 10.1.101.2, xe13, installed 00:30:00, last update 00:30:00 ago
O      1.0.0.2/32 [110/3] via 10.1.101.2, xe13, installed 00:25:48, last update 00:25:48 ago
C      1.1.1.1/32 is directly connected, lo, installed 00:37:24, last update 00:37:24 ago
O      2.2.2.2/32 [110/4] via 10.1.101.2, xe13, installed 00:21:35, last update 00:21:35 ago
O      10.1.12.0/24 [110/2] via 10.1.101.2, xe13, installed 00:30:00, last update 00:30:00 ago
C      10.1.101.0/24 is directly connected, xe13, installed 00:33:18, last update 00:33:18 ago
O      10.1.201.0/24 [110/3] via 10.1.101.2, xe13, installed 00:22:22, last update 00:22:22 ago
C      100.100.1.0/24 is directly connected, xe10.100, installed 00:33:18, last update 00:33:18 ago
C      127.0.0.0/8 is directly connected, lo, installed 00:55:01, last update 00:55:01 ago
B      200.200.1.0/24 [200/0] via 2.2.2.2 (recursive via 10.1.101.2), installed 00:21:29, last update 00:21:29 ago
```

Gateway of last resort is not set

```
PE2#show ip route
```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
 O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
 ia - IS-IS inter area, E - EVPN,
 v - vrf leaked
 * - candidate default

IP Route Table for VRF "default"

```
O      1.0.0.1/32 [110/3] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
O      1.0.0.2/32 [110/2] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
O      1.1.1.1/32 [110/4] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
```

```

C      2.2.2.2/32 is directly connected, lo, installed 00:31:30, last update 00:31:30 ago
O      10.1.12.0/24 [110/2] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44 ago
O      10.1.101.0/24 [110/3] via 10.1.201.2, ce0/1, installed 00:30:44, last update 00:30:44
ago
C      10.1.201.0/24 is directly connected, ce0/1, installed 00:31:30, last update 00:31:30 ago
B      100.100.1.0/24 [200/0] via 1.1.1.1 (recursive via 10.1.201.2), installed 00:30:39, last
update 00:30:39 ago
C      127.0.0.0/8 is directly connected, lo, installed 00:59:47, last update 00:59:47 ago
C      200.200.1.0/24 is directly connected, xe47.100, installed 00:31:30, last update 00:31:30
ago

```

Gateway of last resort is not set

Verify MPLS Labels (ILM or FTN) for SR-MPLS

- ILM table shows incoming label entries learned for SR-OSPF.
- FTN table shows outgoing labels and interfaces used to forward labeled traffic.

```
PE1#show mpls ilm-table
```

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI
P - SR Policy, U - unknown, UPStr - upstream

```

ILM-ECMP LDP: Disabled, SR: Enabled

Code	FEC/VRF/L2CKT	ILM-ID	In-Label	Out-Label	In-Intf	Out-	
Intf/VRF	NextHop		pri	Algo-Num	UpTime	UPStr	peers
O>	2.2.2.2/32	5	16222	16222	N/A	xe13	10.1.101.2
	Yes 0	00:21:50					
O>	1.0.0.2/32	4	16102	16102	N/A	xe13	10.1.101.2
	Yes 0	00:26:03					
O>	1.0.0.1/32	3	16101	16101	N/A	xe13	10.1.101.2
	Yes 0	00:30:15					
O>	1.1.1.1/32	1	16111	NoLabel	N/A	N/A	127.0.0.1
	Yes 0	00:06:29					
O>	10.1.101.2/32	2	24960	3	N/A	xe13	10.1.101.2
	Yes 0	00:30:17					

```
PE1#show mpls forwarding-table
```

```

Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
(m) - FTN mapped over multipath transport, (e) - FTN is ECMP

```

FTN-ECMP LDP: Disabled, SR: Enabled

Code	FEC	FTN-ID	Nhlfe-ID	Tunnel-ID	Pri	Out-Label	Out-	
Intf	ELC	NextHop	Algo-Num	UpTime				
O>	1.0.0.1/32	1	6	-	-	-	-	-
	-	0	00:30:19					
			5	0	Yes	16101	xe13	No
10.1.101.2	-	-						
O>	1.0.0.2/32	2	9	-	-	-	-	-
	-	0	00:26:07					
			8	0	Yes	16102	xe13	No
10.1.101.2	-	-						
O>	2.2.2.2/32	3	12	-	-	-	-	-
	-	0	00:21:54					
			11	0	Yes	16222	xe13	No
10.1.101.2	-	-						

```
PE2#show mpls ilm-table
```

```

Codes: > - installed ILM, * - selected ILM, p - stale ILM, ! - using backup
K - CLI ILM, T - MPLS-TP, s - Stitched ILM
S - SNMP, L - LDP, R - RSVP, C - CRLDP
B - BGP, K - CLI, V - LDP_VC, I - IGP_SHORTCUT
O - OSPF/OSPF6 SR, i - ISIS SR, k - SR CLI

```

```

P - SR Policy,          U - unknown, UPStr - upstream

ILM-ECMP LDP: Disabled, SR: Enabled
Code   FEC/VRF/L2CKT   ILM-ID   In-Label   Out-Label   In-Intf   Out-
Intf/VRF   Nexthop      pri  Algo-Num  UpTime      UPStr peers
O>  2.2.2.2/32        1        16222      Nolabel     N/A       N/A              127.0.0.1
    Yes 0            00:31:53
O>  1.0.0.2/32        4        16102      16102       N/A       ce0/1            10.1.201.2
    Yes 0            00:31:08
O>  1.0.0.1/32        3        16101      16101       N/A       ce0/1            10.1.201.2
    Yes 0            00:31:08
O>  1.1.1.1/32        5        16111      16111       N/A       ce0/1            10.1.201.2
    Yes 0            00:15:45
O>  10.1.201.2/32     2        24960      3           N/A       ce0/1            10.1.201.2
    Yes 0            00:31:46

PE2#show mpls forwarding-table
Codes: > - installed FTN, * - selected FTN, p - stale FTN, ! - using backup
       B - BGP FTN, K - CLI FTN, (t) - tunnel, P - SR Policy FTN, (b) - bypass,
       L - LDP FTN, R - RSVP-TE FTN, S - SNMP FTN, I - IGP-Shortcut,
       U - unknown FTN, O - SR-OSPF FTN, i - SR-ISIS FTN, k - SR-CLI FTN
       (m) - FTN mapped over multipath transport, (e) - FTN is ECMP

FTN-ECMP LDP: Disabled, SR: Enabled
Code   FEC           FTN-ID   Nhlfe-ID   Tunnel-ID   Pri   Out-Label   Out-
Intf   ELC           Nexthop   Algo-Num   UpTime
O>  1.0.0.1/32     1        4          -           -      -          -
    -              0        00:31:11   3          0        Yes  16101      ce0/1      No
10.1.201.2        -        -          -           -      -          -          -
O>  1.0.0.2/32     2        7          -           -      -          -          -
    -              0        00:31:11   6          0        Yes  16102      ce0/1      No
10.1.201.2        -        -          -           -      -          -          -
O>  1.1.1.1/32     3        16         -           -      -          -          -
    -              0        00:15:48   15         0        Yes  16111      ce0/1      No
10.1.201.2        -        -          -           -      -          -          -

```

Verify MPLS Recursive Route Resolution on PE Routers

The `show ip route database` and `show ip route mpls` outputs confirm that PE routers resolve BGP routes via SR-MPLS transport.

On PE1

- The BGP route 200.200.1.0/24 (remote network behind PE2) is learned from the BGP peer 2.2.2.2 (PE2 loopback).
- The BGP next-hop 2.2.2.2 is recursively reachable through OSPF via 10.1.101.2 (toward P1).
- The additional line (recursive via MPLS FTN-ID 3, nhlfe-ix 12) confirms that the route is installed using an MPLS forwarding entry, proving SR-MPLS label forwarding through the core.

```

PE1#show ip route database
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,
       ia - IS-IS inter area, E - EVPN,
       v - vrf leaked
       > - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"
O   *>  1.0.0.1/32 [110/2] via 10.1.101.2, xe13, installed 00:30:04, last update 00:30:04 ago

```

```

O   *> 1.0.0.2/32 [110/3] via 10.1.101.2, xe13, installed 00:25:52, last update 00:25:52 ago
C   *> 1.1.1.1/32 is directly connected, lo, installed 00:37:28, last update 00:37:28 ago
O   1.1.1.1/32 [110/1] is directly connected, lo, installed 00:32:55, last update 00:32:55 ago
O   *> 2.2.2.2/32 [110/4] via 10.1.101.2, xe13, installed 00:21:39, last update 00:21:39 ago
B   2.2.2.2/32 [200/0] via 2.2.2.2 inactive, installed 00:21:33, last update 00:21:33 ago
O   *> 10.1.12.0/24 [110/2] via 10.1.101.2, xe13, installed 00:30:04, last update 00:30:04 ago
C   *> 10.1.101.0/24 is directly connected, xe13, installed 00:33:22, last update 00:33:22 ago
O   10.1.101.0/24 [110/1] is directly connected, xe13, installed 00:32:55, last update 00:32:55
ago
O   *> 10.1.201.0/24 [110/3] via 10.1.101.2, xe13, installed 00:22:26, last update 00:22:26 ago
B   10.1.201.0/24 [200/0] via 2.2.2.2 (recursive via 10.1.101.2), installed 00:21:33, last
update 00:21:33 ago
C   *> 100.100.1.0/24 is directly connected, xe10.100, installed 00:33:22, last update 00:33:22
ago
C   *> 127.0.0.0/8 is directly connected, lo, installed 00:55:05, last update 00:55:05 ago
B   *> 200.200.1.0/24 [200/0] via 2.2.2.2 (recursive via 10.1.101.2), installed 00:21:33, last
update 00:21:33 ago
    *> [200/0] via 2.2.2.2 (recursive via MPLS FTN-ID 3, nhlfe-ix 12)

```

Total number of IPv4 routes 15

Gateway of last resort is not set

PE1#show ip route mpls

IP Route Table for VRF "default"

```

B   200.200.1.0/24 [200/0] via 2.2.2.2 (recursive via MPLS FTN-ID 3, nhlfe-ix 12), installed
00:21:38, last update 00:21:38 ago

```

Gateway of last resort is not set

On PE2

- The BGP route 100.100.1.0/24 (remote network behind PE1) is learned from the BGP peer 1.1.1.1 (PE1 loopback).
- The BGP next-hop 1.1.1.1 is recursively reachable through OSPF via 10.1.201.2 (toward P2).
- The additional line (recursive via MPLS FTN-ID 3, nhlfe-ix 16) confirms MPLS forwarding entry selection for SR-MPLS transport.

PE2#show ip route database

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP

O - OSPF, IA - OSPF inter area

N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2

E1 - OSPF external type 1, E2 - OSPF external type 2

i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2,

ia - IS-IS inter area, E - EVPN,

v - vrf leaked

> - selected route, * - FIB route, p - stale info

IP Route Table for VRF "default"

```

O   *> 1.0.0.1/32 [110/3] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
O   *> 1.0.0.2/32 [110/2] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
O   *> 1.1.1.1/32 [110/4] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
B   1.1.1.1/32 [200/0] via 1.1.1.1 inactive, installed 00:30:51, last update 00:30:51 ago
C   *> 2.2.2.2/32 is directly connected, lo, installed 00:31:42, last update 00:31:42 ago
O   2.2.2.2/32 [110/1] is directly connected, lo, installed 00:31:42, last update 00:31:42 ago
O   *> 10.1.12.0/24 [110/2] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
O   *> 10.1.101.0/24 [110/3] via 10.1.201.2, ce0/1, installed 00:30:56, last update 00:30:56 ago
B   10.1.101.0/24 [200/0] via 1.1.1.1 (recursive via 10.1.201.2), installed 00:30:51, last
update 00:30:51 ago
C   *> 10.1.201.0/24 is directly connected, ce0/1, installed 00:31:42, last update 00:31:42 ago
O   10.1.201.0/24 [110/1] is directly connected, ce0/1, installed 00:31:42, last update
00:31:42 ago
B   *> 100.100.1.0/24 [200/0] via 1.1.1.1 (recursive via 10.1.201.2), installed 00:30:51, last
update 00:30:51 ago
    *> [200/0] via 1.1.1.1 (recursive via MPLS FTN-ID 3, nhlfe-ix 16)

```

```

C    *> 127.0.0.0/8 is directly connected, lo, installed 00:59:59, last update 00:59:59 ago
C    *> 200.200.1.0/24 is directly connected, xe47.100, installed 00:31:42, last update 00:31:42
ago

Total number of IPv4 routes 15

Gateway of last resort is not set

PE2#show ip route mpls
IP Route Table for VRF "default"
B      100.100.1.0/24 [200/0] via 1.1.1.1 (recursive via MPLS FTN-ID 3, nhife-ix 16), installed
00:30:58, last update 00:30:58 ago

Gateway of last resort is not set

```

Test End-to-End Reachability (SR-MPLS)

Ping or traceroute must reach the remote customer subnet through SR-MPLS transport (labels are imposed at ingress and forwarded across the core).

From PE1 to CE2

```

PE1#ping 200.200.1.1
Press CTRL+C to exit
PING 200.200.1.1 (200.200.1.1) 100(128) bytes of data.
108 bytes from 200.200.1.1: icmp_seq=1 ttl=64 time=0.715 ms
108 bytes from 200.200.1.1: icmp_seq=2 ttl=64 time=0.665 ms
108 bytes from 200.200.1.1: icmp_seq=3 ttl=64 time=0.602 ms

--- 200.200.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time 2073ms
rtt min/avg/max/mdev = 0.602/0.660/0.715/0.046 ms

PE1#traceroute 200.200.1.1
traceroute to 200.200.1.1 (200.200.1.1), 30 hops max, 60 byte packets
 1 200.200.1.1 (200.200.1.1) 0.898 ms 0.874 ms 0.901 ms

```

From PE2 to CE1

```

PE2#ping 100.100.1.1
Press CTRL+C to exit
PING 100.100.1.1 (100.100.1.1) 100(128) bytes of data.
108 bytes from 100.100.1.1: icmp_seq=1 ttl=64 time=0.677 ms
108 bytes from 100.100.1.1: icmp_seq=2 ttl=64 time=0.642 ms
108 bytes from 100.100.1.1: icmp_seq=3 ttl=64 time=0.595 ms
108 bytes from 100.100.1.1: icmp_seq=4 ttl=64 time=0.617 ms

--- 100.100.1.1 ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3104ms
rtt min/avg/max/mdev = 0.595/0.632/0.677/0.030 ms

PE2#traceroute 100.100.1.1
traceroute to 100.100.1.1 (100.100.1.1), 30 hops max, 60 byte packets
 1 100.100.1.1 (100.100.1.1) 0.782 ms 0.881 ms 0.915 ms

```

Observed Behavior (SR-MPLS)

- **Segment Routing State:** PE1 and PE2 report SR_ENABLED under OSPF, confirming SR-MPLS is active and operational in the transport domain.
- **Prefix-SID Distribution:** PE1 (1.1.1.1/32, SID 111), P1 (1.0.0.1/32, SID 101), P2 (1.0.0.2/32, SID 102), and PE2 (2.2.2.2/32, SID 222) are present in the SR mapping table, confirming end-to-end Prefix-SID learning

across the core.

- **MPLS Forwarding (SR Labels):** SR-OSPF label entries are installed on **PE1** and **PE2**, and labels are used for transport forwarding across **P1** and **P2** (no LDP sessions are required).
- **BGP and MPLS Interaction:** **PE1** learns the remote customer prefix **200.200.1.0/24** from **PE2 (2.2.2.2)** and resolves it through SR-MPLS transport (`recursive via MPLS FTN-ID 3, nhlfe-ix 12`). Similarly, **PE2** learns **100.100.1.0/24** from **PE1 (1.1.1.1)** and resolves it through SR-MPLS (`recursive via MPLS FTN-ID 3, nhlfe-ix 16`).
- **Validation confirms:**
 - IP routes are correct (**PE1/PE2** `show ip route`).
 - SR state and Prefix-SIDs are present (**PE1/PE2** `show ip ospf segment-routing state and mapping-table active`).
 - MPLS SR label tables are populated (**PE1/PE2** `show mpls ilm-table` and `show mpls forwarding-table`).
 - End-to-end connectivity works (`ping` and `traceroute`).

Implementation Example

Use Case: Simplified MPLS Core for BGP-Free Operation

In a large network backbone, the service provider's core is designed for label-switched forwarding using MPLS. Dynamic routing, such as BGP, is enabled only at the network edges where service routes are originated or terminated. The core routers focus solely on MPLS transport, without maintaining full BGP tables.

With Recursive Next-Hop Resolution, OcNOS edge or aggregation routers can automatically resolve BGP or static routes through available MPLS paths, even when the next-hop is not directly reachable via IP. This allows traffic to traverse the MPLS core seamlessly, enabling a scalable and operationally efficient BGP-free core design. The feature also leverages MPLS protection mechanisms for rapid failover and improved resiliency within the core.

Revised Commands

show ip route mpls

Added the `mpls` and `mpls summary` parameters to the `show ip route` command to display routing table entries resolved through MPLS next-hops (NHLFE). For more details, refer to the `show ip route` command section in the *OcNOS System Management Guide*.

show ip route database

Added the `MPLS FTN-ID <ID>`, `nhlfe-ix <INDEX>` field to the `show ip route database` command output. This field specifies the MPLS next-hop used for recursive resolution. For more details, refer to the `show ip route database` command section in the *OcNOS System Management Guide*.

show rib txlist

The `show rib txlist` ([page 331](#)) command includes the fields (`req_mpls`, `new_req_mpls`, `fib_notify_mpls`, `seq_mpls`) to display MPLS-specific route installation and synchronization details between RIB and FIB.

Troubleshooting

The following steps can be used to validate recursive next-hop resolution through MPLS paths:

1. Verify that RIB has MPLS transport information: Use the diagnostic command `show rib mpls ipv4` (hidden debug command) to confirm that MPLS transport details are received and available for resolution.
2. Verify that recursive routes are installed: Check the routing table entries resolved via MPLS next-hops using the show commands `show ip route (|summary | database)` and `show ip route mpls (|summary)`.
3. Verify recursive route installation in hardware: Confirm that MPLS-resolved prefixes are programmed in the forwarding tables using the following show commands:
 - `show hsl prefix-table ipv4 fib-id 0`
 - For MPLS transport validation:
 - LDP or SR-MPLS tunnels: `show hsl mpls ftn nhlfe-ix <index>`
 - RSVP-TE tunnels: `show hsl mpls tunnel tunnel-id <id>`

show rib txlist

Use this command to display internal Routing Information Base (RIB) transaction lists that are pending or currently being transmitted to the Forwarding Information Base (FIB). This command helps monitor routes awaiting installation or synchronization between the routing and forwarding planes. It is mainly used for debugging and operational verification of route processing states.

Command Syntax

```
show rib (txlist | pending-txlist) (ipv4 | ipv6) (vrf WORD|)
```

Parameters

txlist

Displays the list of routes that are currently being transmitted from RIB to FIB.

pending-txlist

Displays the list of routes queued for transmission to the FIB.

ipv4

Displays information for IPv4 routes.

ipv6

Displays information for IPv6 routes.

vrf WORD

Specifies the VRF name for which the RIB transaction information is displayed.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3.

Added the fields (req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls) in OcNOS version 7.0.0.

Example

The fields (req_mpls, new_req_mpls, fib_notify_mpls, seq_mpls) show MPLS-specific route installation and synchronization details between RIB and FIB.

```
#show rib txlist ipv4
IPv4:
0x02739a6c: 127.0.0.0/8      req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
0x02810cdc: 1.2.0.0/24      req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(5)
0x0281071c: 1.1.1.1/32      req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(5)
0x028d9c5c: 2.2.2.2/32      req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(4)
0x028da42c: 11.11.11.11/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
0x028dab2c: 12.12.12.12/32     req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
```

```

0x028daf7c: 13.13.13.13/32      req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
0x028db0cc: 9.9.9.9/32         req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
0x028db32c: 10.10.10.10/32      req(0) new_req(0) fib_notify(0) seq(0) req_mpls(0) new_req_mpls(0)
fib_notify_mpls(0) seq_mpls(0) lock(3)
MARKER MPLS*
MARKER*

#show rib pending-txlist ipv4 vrf VRF3
IPv4 pending:
0x02634680: 192.0.5.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634830: 192.0.6.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x026349e0: 192.0.7.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02634bb0: 192.0.8.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652490: 192.0.9.0/24      req(1) new_req(0) fib_notify(0) seq(0) lock(3)
0x02652660: 192.0.10.0/24     req(1) new_req(0) fib_notify(0) seq(0) lock(3)

```

Glossary

The following provides definitions for key terms or abbreviations and their meanings used throughout this document:

Key Term or Acronym	Description
BGP-Free Core	Network design where MPLS transport is used in the core without running BGP between core routers. Edge devices handle route advertisement and resolution.
Routing Information Base (RIB)	The routing table in control plane that maintains best route selections and recursive next-hop information.
Forwarding Information Base (FIB)	The hardware forwarding table that programs routes installed by the RIB for packet forwarding.
Recursive Resolution	The process of resolving a route's next-hop through another route or transport (e.g., MPLS tunnel) instead of a directly connected IP interface.
Multiprotocol Label Switching (MPLS)	Label-based forwarding technology used for scalable and efficient transport within the core network.
Next Hop Label Forwarding Entry (NHLFE)	MPLS data structure that defines label operations (push, swap, pop) and the next-hop interface for labeled forwarding.
FEC-to-NHLFE Mapping (FTN)	Mapping between Forwarding Equivalence Class (FEC) and NHLFE entries used to associate IP prefixes with MPLS label paths.
Label Distribution Protocol (LDP)	Protocol that distributes labels for MPLS forwarding paths dynamically across routers.
Resource Reservation Protocol - Traffic Engineering (RSVP-TE)	Protocol used to establish MPLS-TE LSPs with defined resource and path constraints.
Segment Routing - MPLS (SR-MPLS)	MPLS forwarding model where segment identifiers (SIDs) define packet paths without per-flow signaling.
Equal-Cost Multi-Path (ECMP)	Mechanism that enables load sharing across multiple equal-cost routes.
Fast Reroute (FRR)	Protection mechanism that provides rapid failover by pre-installing backup label paths.
Topology Independent Loop-	Advanced FRR mechanism ensuring near-instantaneous recovery for link or

Key Term or Acronym	Description
Free Alternate (TI-LFA)	node failures.
Hardware Support Layer (HSL)	OcNOS subsystem responsible for programming and managing hardware forwarding entries in FIB.
FTN-ID	Identifier assigned to an MPLS Forwarding Equivalence Class entry used in recursive resolution.
NHLFE-IX	Index that uniquely identifies the NHLFE used by the MPLS forwarding path.
BGP Multipath	Capability that allows installation of multiple equal-cost BGP paths; not supported in current release for BGP-Free Core scenarios.