



OcNOS®

Open Compute Network Operating System for Data Centers

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
Routing and Security Enhancements	12
BGP RPKI-Based Route Validation	13
Overview	13
Feature Characteristics	13
Benefits	13
Configuration	14
Prerequisites	14
Topology	14
Validation	16
Implementation Examples	18
Commands	18
bgp rpk server	19
bgp origin-as validation-enable	21
bgp origin-as bestpath	22
match rpk	23
show bgp rpk table ipv4	24
show bgp rpk table ipv6	26
show bgp origin-as validity ipv4	28
show bgp origin-as validity ipv6	29
show bgp rpk server	31
Glossary	32
EVPN AF Route-Maps and Route Filtering	34
Overview	34
Feature Characteristics	34
Benefits	34
Prerequisites	34

Limitation	35
Configuration	35
Validation	37
Implementation Examples	44
Commands	44
mac-list	45
match mac address list	46
match evpn-route-type	47
show mac-list	48
Glossary	48
BGP Labeled Unicast Next Hop Self in Route-Map	49
Overview	49
Feature Characteristics	49
Limitations	49
Prerequisites	49
Configuration	49
Configuration Snapshot	51
Commands	109
set ip next-hop self	109
EVPN L3 Gateway with VXLAN Stitching	110
HPC or Artificial Intelligence Networking	111
Dynamically Adjusts Explicit Congestion Notification Marking Threshold Values	113
PFC Deadlock Detection and Recovery	114
Overview	114
Priority Flow Control Pause Frames	114
Workflow of PFC Frames	114
PFC deadlock:	114
Feature Characteristics	115
Deadlock Detection	115
Deadlock Recovery	115
Benefits	115
Prerequisites	116
Configuration	116
Topology	116
Timer mode	116
PFC state XON mode	117
Global Mode	117
Validation	117
Using manual recovery on an interface	118
Validation	118
Spine and Leaf Topology	119
Configure Leaf 1, Leaf 2, Leaf 3, Leaf 4, Spine 1 and Spine 2	120
Validation	137
PFC DD Commands	139
clear priority-flow-control deadlock-status	140
priority-flow-control deadlock manual-recovery	141

priority-flow-control deadlock recovery-action drop	142
priority-flow-control deadlock recovery-mode timer	143
priority-flow-control deadlock recovery-mode pfc-state-xon	144
show priority-flow-control deadlock-status	145
Implementation Examples	145
Glossary	146
PFC Deadlock Detection and Recovery	146
Overview	146
Feature Characteristics	147
Benefits	148
Prerequisites	148
Configuring PFC Deadlock Detection and Recovery	148
Topology	148
Configure Using Timer mode	149
PFC State XON mode	152
Global Mode	153
Manual Recovery	153
Clear PFC Deadlock	154
CLI Commands	156
clear priority-flow-control deadlock-status	156
priority-flow-control deadlock manual-recovery	157
priority-flow-control deadlock recovery-action drop	157
priority-flow-control deadlock recovery-mode timer	158
priority-flow-control deadlock recovery-mode pfc-state-xon	159
show priority-flow-control deadlock-status	159
Implementation Examples	160
Glossary	161
PFC Frames and ECN Packets Monitoring	162
Overview	162
Feature Characteristics	162
Benefits	162
Prerequisites	162
Configuring PFC Frames and ECN Packets Monitoring	162
Topology	163
Configuration for ECN Marking and PFC Pausing	163
Sample Show Running Configuration on Switch	166
Validation	170
ECN Validation	170
PFC Validation	171
PFC-ECN Commands	175
monitor ecn	176
monitor pfc	176
Glossary	177
PFC Frames and ECN Packets Monitoring	177
Overview	177
Feature Characteristics	177
Benefits	177

Prerequisites	178
CLI Commands	178
monitor ecn	178
monitor pfc	178
Glossary	179
ECN and PFC Support for Lossless VxLAN Transport	180
Switch Packet Buffer Tuning	181
Network Management and Automation	182
Mirror Filtered Data to CPU	183
NetConf Access Control Model User Guide	184
Overview	184
Feature Characteristics	184
Configuration	185
Initial NACM Configuration	185
Prerequisites	186
Common NACM Rule Fields	186
Creating NetConf RPC for NACM	189
RPC Configurations for NACM	189
Glossary	198
sFlow - Sample Packet Monitoring for Multiple Interfaces	199
Overview	199
Features Characteristics	200
Benefits	200
Prerequisites	200
Configuration	200
Topology	200
Validation	203
Configuring sFlow with User Defined VRFs	203
Implementation Examples	206
Commands	212
no sflow collector-id	213
sflow collector	214
Troubleshooting	215
Glossary	215
VxLAN OAM for Overlay Networks	216
CLI-Script and CLI-Shell	217
Overview	217
Feature Characteristics	217
Benefits	217
Limitations	217
Configuration	218
CLI-Script Configuration	218
CLI-Shell Configuration	218
Configuration for Delay and Message Commands	218
Validation	219
Configuration Snapshot	219

Implementation Examples	220
CLI-Script and CLI-Shell Commands	221
cli-script	222
cli-script line command	223
cli-script-end	224
show cli-script	225
load-cli-script	226
exec-shell	227
delay	228
message	229
show running-config extended	230
show cli-script content all	231
no cli-script	232
copy running-config-ext <remote-location>	233
System Limits and Counters	234
Overview	234
Feature Characteristics	234
Benefits	234
System Limits and Counters Limitation	235
Data Characteristics	235
Security and Access	235
Dependencies	235
Platform-specific	235
Encoding	235
System Limits and Counters Configuration	235
Topology	235
Use Case: Verify Total Number of IPv4 Routes Installed	236
System Limits and Counters Implementation Example	242
Scenario 1: Resource Audit During Large-Scale Migration	242
Scenario 2: Monitoring via Network Management Systems	242
System Limits and Counters Commands	242
System Limits and Counters Revised Commands	242
show access-lists	243
show ip vrf	244
show ip route	245
show ipv6 route	247
show interface	249
System Limit Counters Troubleshooting	252
Show Output Displays Blank or Partial Results	252
gNMI Returns “Unsupported Encoding” or Missing Fields	253
API Retrieval Fails for Specific Resource Paths	253
System Limit Counters Glossary	253
Layer 2 or Layer 3 Overlay Networking	255
Layer 3 Sub-interface	256
Overview	256
Feature Characteristics	256

Benefits	256
Limitations	256
Configuration	257
Topology	257
Creating a Sub-interface	257
Creating a Sub-interface with Encapsulation	257
Validation	259
show interface brief	259
show ip interface brief	259
show ip ospf neighbor with VRF enabled	259
show ip route with VRF enabled	260
Layer 3 Sub-interface Commands	262
encapsulation	263
interface IFNAME.SUBINTERFACE_ID	264
show interface IFNAME.SUBINTERFACE_ID	265
Implementation Examples	266
Troubleshooting	267
Glossary	268

PREFACE

About this Guide

This guide describes how to configure Key Features in OcNOS.

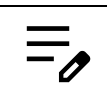
Audience

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

Conventions

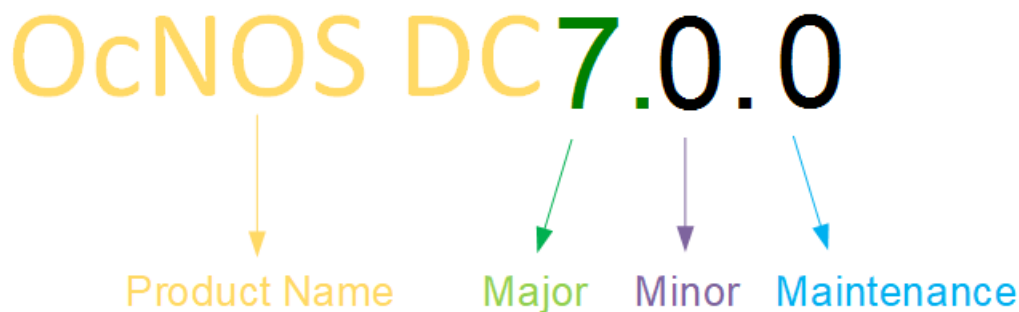
The [Table 1](#) table shows the conventions used in this guide.

Table 1. Conventions

Convention	Description
<i>Italics</i>	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.

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 13
 - Overview 13
 - Feature Characteristics 13
 - Benefits 13
 - Configuration 14
 - Implementation Examples 18
 - Commands 18
 - Glossary 32
- EVPN AF Route-Maps and Route Filtering 34
 - Overview 34
 - Feature Characteristics 34
 - Benefits 34
 - Prerequisites 34
 - Configuration 35
 - Implementation Examples 44
 - Commands 44
 - Glossary 48
- BGP Labeled Unicast Next Hop Self in Route-Map 49
 - Overview 49
 - Feature Characteristics 49
 - Limitations 49
 - Prerequisites 49
 - Configuration 49
 - Commands 109
- EVPN L3 Gateway with VXLAN Stitching 110

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.

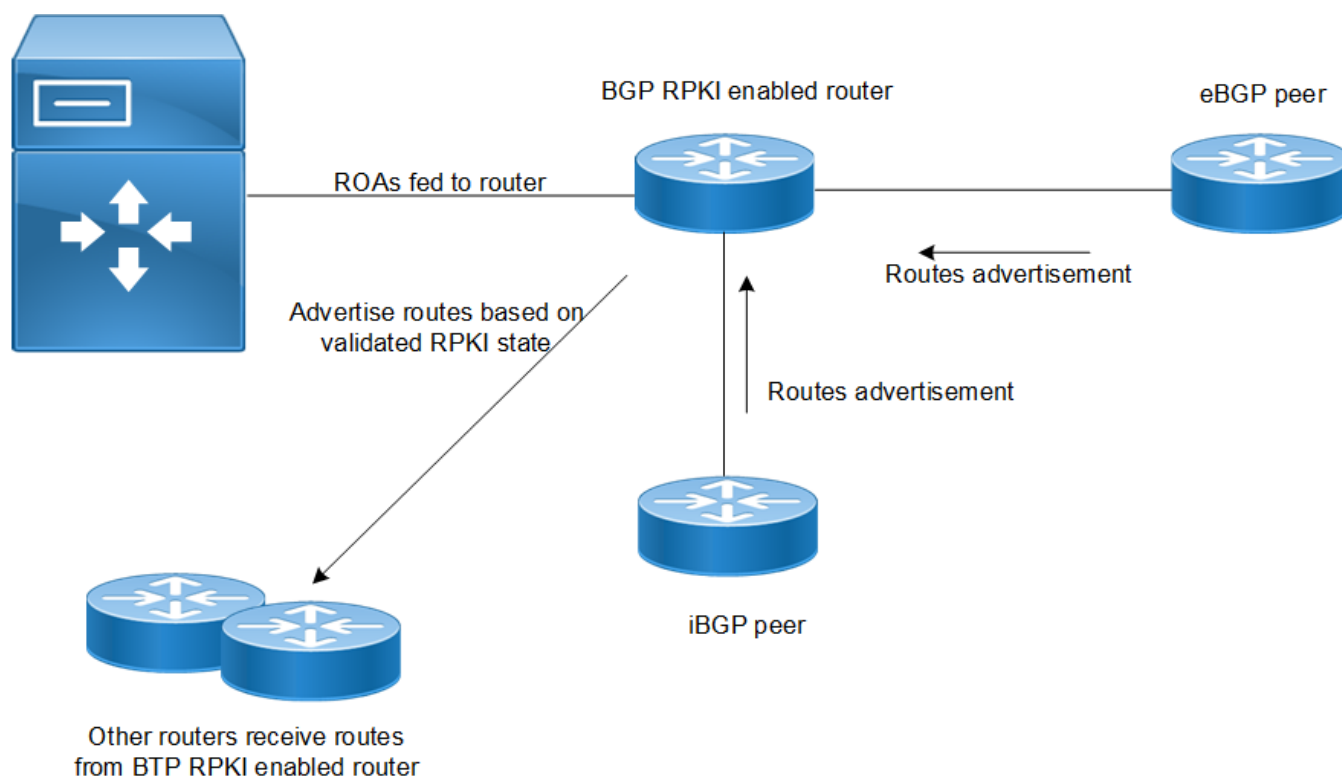
Configuration

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*	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.

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 1. 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 \(page 34\)](#).

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 65000 0 0	4	1302	15	0	0	04:59:49	5	1	3 1

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

Hop	Network	Next Metric	Next 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 Type	VTEP- Status	MAC move	AccessPortDesc
205	irb305	----	----	0011.2233.4455 Static Local	-----	0	-----
205	-----	----	----	0011.2233.4567 Static Remote	-----	0	-----
205	irb305	----	----	6821.5f1f.5220 Static Local	-----	0	-----
205	-----	----	----	e8c5.7a69.45ed 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
Nextthop : 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
Nextthop : 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
Nextthop : 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
Nextthop : 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
Nextthop : 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
Nextthop : 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
Nextthop : 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
Nextthop : 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
Nextthop : 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 l [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 l [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 l [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 2. 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
11/13/2025 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 l2-circuit PE2-To-PE4-1297 4294961297 192.168.36.4
  control-word
!
mpls l2-circuit PE2-To-PE4-1298 4294961298 192.168.36.4
  control-word
!
mpls l2-circuit PE2-To-PE4-1299 4294961299 192.168.36.4
  control-word
!
mpls l2-circuit PE2-To-PE4-1300 4294961300 192.168.36.4
  control-word
!
mpls l2-circuit vc2000 4294967290 192.168.36.11
!
mpls l2-circuit vc2001 4294967290 192.168.36.12
!
mpls l2-circuit vc2004 4294967291 192.168.36.11
!
mpls l2-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-IRBVRF3003-IPv6
seq 5 permit ::/0
!
ipv6 prefix-list PL-DNS-SERVERS-IRBVRF3003-IPv6
seq 5 permit 2001:db8:3003:53::53/128
!
ipv6 prefix-list PL-NEXTHOP-IRBVRF3003-IPv6
seq 5 permit 80:12:1::254/128
seq 10 permit 2001:bd8:103::1/128
!
ipv6 prefix-list PL-NOMETRIC-IRBVRF3003-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-IRBVR3003-IPv4 permit 1001  
  match ip address prefix-list PL-NOMETRIC-IRBVR3003-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-IRBVR3003-IPv6 permit 1002  
  match ip address prefix-list PL-CUST-SUBNETS-IRBVR3003-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-IRBVR3003-IPv6 permit 1003  
  match ip address prefix-list PL-DENY-DEFAULT-IRBVR3003-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-IRBVR3003-IPv6 permit 1004  
  match ip address prefix-list PL-DNS-SERVERS-IRBVR3003-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-IRBVR3003-IPv6 permit 1005
  match ip next-hop prefix-list PL-NEXTHOP-IRBVR3003-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-IRBVR3003-IPv6 permit 1006
  match community CUSTOMER_ROUTES_3015_IRBVR3003_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-IRBVR3003-IPv6 permit 1007
  match extcommunity CUSTOMER_EXTENDED_COMM_3016_IRBVR3003_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-IRBVR3003-IPv6 permit 1008
  match extcommunity CUSTOMER_EXTENDED_COMM_3017_4byte_IRBVR3003_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-IRBVR3003-IPv6 permit 1009
  match large-community CUSTOMER_LARGE_COMM_IRBVR3003_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-IRBVR3003-IPv6 permit 1010
  match as-path ASPATH-IRBVR3003-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 xe1.113
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 xe1.114
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 xe1.115
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 xe1.116
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 xe1.117
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 xe1.118
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 xe1.119
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 xe1.120
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 xe1.300  
!  
interface xe1.890  
  description for ### IPv4 eBGP  
  encapsulation dot1q 890  
  load-interval 30  
  ip address 190.160.2.1/24  
  mtu 9216  
!  
interface xe1.891  
  description for ### 6PE  
  encapsulation dot1q 891  
  load-interval 30  
  ipv6 address 3601::1/64  
  mtu 9216  
!  
interface xe1.892  
  description for ### 6PE  
  encapsulation dot1q 892  
  load-interval 30  
  ipv6 address 3602::1/64  
  mtu 9216  
!  
interface xe1.893  
  description for ### 6PE  
  encapsulation dot1q 893  
  load-interval 30  
  ipv6 address 3603::1/64  
  mtu 9216  
!  
interface xe1.894  
  description for ### 6PE  
  encapsulation dot1q 894  
  load-interval 30  
  ipv6 address 3604::1/64  
  mtu 9216  
!  
interface xe1.895  
  description for ### 6PE  
  encapsulation dot1q 895  
  load-interval 30  
  ipv6 address 3605::1/64  
  mtu 9216  
!  
interface xe1.896  
  description for ### 6PE  
  encapsulation dot1q 896  
  load-interval 30  
  ipv6 address 3606::1/64  
  mtu 9216  
!  
interface xe1.897  
  description for ### 6PE  
  encapsulation dot1q 897  
  load-interval 30  
  ipv6 address 3607::1/64  
  mtu 9216  
!  
interface xe1.898  
  description for ### 6PE  
  encapsulation dot1q 898  
  load-interval 30  
  ipv6 address 3608::1/64  
  mtu 9216  
!  
interface xe1.899  
  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 ### e1ine908
  encapsulation dot1q 908
  load-interval 30
  mtu 9216
  access-if-evpn
    map vpn-id 16772908
!
interface xe1.909 switchport
  description for ### e1ine909
  encapsulation dot1q 909
  load-interval 30
  mtu 9216
  access-if-evpn
    map vpn-id 16772909
!
interface xe1.910 switchport
  description for ### e1ine910
  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 vpv4 all-vrfs per-vrf
mpls label mode vpv6 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 0xffff87e79fdacd4e7
  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      39      3     0     0 00:12:04      1
3.3.3.3      4 4200000001   35      43      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   Nhlfe-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
      5          -
      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     N/A
      Yes N/A      00:20:13
B>    VRF1          1        24320     Nolabel    N/A     N/A     N/A
      Yes N/A      00:20:13
B>    2.2.2.2/32    4        24323     26240     N/A     N/A     2.2.2.2
      Yes N/A      00:20:06
B>    3.3.3.3/32    3        24322     26241     N/A     N/A     2.2.2.2
      Yes N/A      00:10:34
B>    1.1.1.1/32    5        24324     Nolabel    N/A     N/A     127.0.0.1
      Yes N/A      00:20:09
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	Nextthop	FTN-ID	Nhlfe-ID	Tunnel-ID	Pri	Out-Label	Out-
Intf	ELC		Algo-Num	UpTime				
B>	1.1.1.1/32	26242	3	11	-		N/A	00:11:11
R								
(t)>	2.2.2.2/32		1	3	5001	Yes	24320	ce0 No 3
0.1.1.2	N/A		00:21:09					
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	Nextthop		pri	Algo-Num	UpTime	UPStr peers
B>	3.3.3.3/32	4	24962	Nolabel	N/A	N/A 127.0.0.1
	Yes N/A	00:21:12				
B>	VRF1	1	24960	Nolabel	N/A	N/A
	Yes N/A	00:21:22				
R>	3.3.3.3/32	3	24320	Nolabel	N/A	N/A 127.0.0.1
	Yes N/A	00:21:18	1			
B>	VRF2	2	24961	Nolabel	N/A	N/A
	Yes N/A	00:21:22				
B>	2.2.2.2/32	5	24963	26240	N/A	N/A 2.2.2.2
	Yes N/A	00:21:06				
B>	1.1.1.1/32	6	24964	26242	N/A	N/A 2.2.2.2
	Yes N/A	00:11:14				

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
..
```

EVPN L3 Gateway with VXLAN Stitching

OcNOS supports Layer 3 EVPN Gateway Stitching, enabling seamless IP connectivity between independent VXLAN EVPN domains or VRFs without merging their control planes or Layer 2 broadcast domains. The feature uses EVPN Type-5 route stitching at border leaf or spine devices to facilitate scalable and secure inter-domain routing across multi-site data centers, multi-PoD fabrics, and hybrid cloud environments.

For more details, refer to the [EVPN L3 Gateway with VXLAN Stitching](#) section in OcNOS Virtual Extensible LAN Guide, Release 7.0.0.

HPC OR ARTIFICIAL INTELLIGENCE NETWORKING

OcNOS provides high-performance networking features designed to meet the demanding bandwidth, latency, and scalability needs of High-Performance Computing (HPC) and AI workloads, ensuring efficient east-west traffic handling and low-latency interconnect performance.

Dynamically Adjusts Explicit Congestion Notification Marking Threshold Values	113
PFC Deadlock Detection and Recovery	114
Overview	114
Feature Characteristics	115
Benefits	115
Prerequisites	116
Configuration	116
Topology	116
Timer mode	116
PFC state XON mode	117
Global Mode	117
Validation	117
Using manual recovery on an interface	118
Spine and Leaf Topology	119
PFC DD Commands	139
Implementation Examples	145
Glossary	146
PFC Deadlock Detection and Recovery	146
Overview	146
Configuring PFC Deadlock Detection and Recovery	148
CLI Commands	156
Implementation Examples	160
Glossary	161
PFC Frames and ECN Packets Monitoring	162
Overview	162
Prerequisites	162
Configuring PFC Frames and ECN Packets Monitoring	162
Validation	170
PFC-ECN Commands	175
Glossary	177
PFC Frames and ECN Packets Monitoring	177
Overview	177

CLI Commands	178
Glossary	179
ECN and PFC Support for Lossless VxLAN Transport	180
Switch Packet Buffer Tuning	181

Dynamically Adjusts Explicit Congestion Notification Marking Threshold Values

OcNOS is enhanced to support lossless Ethernet fabrics for AI/ML workloads through Dynamic Explicit Congestion Notification (D-ECN) method on Broadcom Tomahawk5 platforms.

D-ECN allows users to adjust the ECN thresholds using D-ECN-ON-Offset and D-ECN-OFF-Offset settings that provides capability to enable precise congestion marking based on shared buffer usage.

Unlike traditional methods that depend on packet drops, D-ECN enhances efficiency by marking IP headers to indicate congestion, prompting receivers to signal senders to adjust transmission rates.

For further details, refer to [Dynamic ECN Marking](#) section in the *OcNOS Quality of Service Guide*, Release 7.0.0.

PFC Deadlock Detection and Recovery

Overview

Priority-based Flow Control (PFC) helps manage traffic in networks by pausing specific flows during congestion. However, under certain conditions, a deadlock can occur when cyclical dependencies between flows create a loop of PFC pause events that prevents traffic from making forward progress indefinitely.

Priority Flow Control Pause Frames

PFC uses the standard pause frame mechanism with an additional 14 bytes of padding in the frame. This padding contains a 2-byte value for each of the eight priority classes, specifying the pause time in quanta for that class.

Example for priority pause frame:

Pause Frame:

```
Control Opcode: 0x0101 (Priority Pause)
Pause Time (8 priority classes):
Class 0: 0x0000 (no pause)
Class 1: 0x1234 (pause time in quanta)
Class 2: 0x5678 (pause time in quanta)
...
Class 7: 0x9ABC (pause time in quanta)
```

In the example above, the pause time in quanta field defines if the pause frame has XON or XOFF set for that class:

- **XON (X-On):** A control signal sent from the receiver to the transmitter to indicate readiness to accept data. For example, Class 0 represents a no-pause condition.
- **XOFF (X-Off):** A control signal sent from the receiver to the transmitter indicating that it cannot accept additional data due to congestion. For example, Classes 1, 2, and 7 above specify a non-zero pause time (in quanta), signaling the transmitter to temporarily halt transmission.

Workflow of PFC Frames

- **Transmission:** The transmitter sends data to the receiver.
- **XOFF:** The receiver sends an XOFF signal to the transmitter, indicating that it is congested and cannot process more data.
- **Pause:** The transmitter pauses sending data to the receiver.
- **XON:** The receiver sends an XON signal to the transmitter, indicating that it is ready to receive data again.
- **Resume:** The transmitter resumes sending data to the receiver.

PFC deadlock:

A deadlock may occur when the receiver continuously sends XOFF signals for one or more classes, preventing the transmitter from sending any traffic. This feature is designed to detect such deadlocks and initiate recovery mechanisms.

To handle such critical situation, the OcNOS system provides PFC Deadlock Detection and Recovery capability. This chapter describes how to:

- Enable PFC deadlock detection and recovery on a specific interface using

- [Timer mode \(page 116\)](#)
- [PFC state XON mode \(page 117\)](#)
- Configure the global PFC deadlock detection and recovery action to drop
 - [Global Mode \(page 117\)](#)

Feature Characteristics

Deadlock Detection

The system monitors PFC queues for extended periods in the XOFF state.

If a queue remains paused beyond a configurable threshold, a deadlock event is declared.

An interrupt is raised to inform software of the detected deadlock.

Deadlock Recovery

Once a deadlock is detected, software moves the affected queue into an ignore PFC XOFF state, allowing traffic scheduling to resume.

Recovery can be configured on a per-interface basis and supports three modes:

- **Timer Mode:** Recovery ends automatically after a user-defined time interval. The system then clears the interrupt and restarts the detection timer. It is an automatic recovery method and recovery starts after a configurable detection-multiplier times time-granularity period. During that period, traffic will be allowed by default, but can also be dropped if the configuration priority-flow-control deadlock recovery-action drop is set. Recovery also ends automatically after an optionally configurable recovery-time period.



Note: Traffic will gradually decrease to zero if the recovery-mode timer is not configured; otherwise, it will continue indefinitely.

- **PFC-State-XON Mode:** Recovery ends when the interface receives a PFC XON frame, signaling that the pause condition is lifted.
- **Manual Mode:** Recovery requires explicit user action with CLI commands. This option is only valid if no automatic recovery mode is configured.

Limitation:

Manual recovery mode is not supported in Trident3 (TR3) platforms or Tomahawk3 (TH3) platforms.

Trident3 (TR3) platforms support deadlock recovery only in timer mode.

Trident3 (TR3) platforms do not support 1ms time granularity.

Tomahawk 2 (TH2) series platforms are not supported.

Benefits

Prevents indefinite traffic stalls due to PFC loops.

Provides flexible recovery options (automatic or manual).

Improves network reliability in environments that rely on PFC.

Prerequisites

The device should be enabled with PFC.

Configuration

PFC feature supports deadlock detection and recovery. This chapter shows how to:

1. Enable PFC deadlock detection and recovery on an interface
2. Set global PFC deadlock detection and recovery action to drop

Topology

This topology illustrates a spine-leaf router architecture where Priority Flow Control (PFC) manages traffic at the queue level, optimizing the flow from spine to leaf while minimizing packet loss and enhancing overall quality of service.

Figure 3. PFC Enabled Bridge



Configuring an Interface for PFC Deadlock Detection and Recovery on interface can be done in Timer mode or XON mode.



Note: Refer to "PFC Deadlock Detection and Recovery" section in Layer 2 Configuration guide for EVPN-VxLAN topology configuration.

Timer mode

Execute the following steps to configure PFC on both interfaces on leaf router.

1. Set the IP address.

```
(config-if)#ip address 1.1.1.1/24
```

2. Enable the PFC. Configure the advertise flag and start sending DCBX TLVs in LLDP messages.

```
(config-if)#priority-flow-control mode on
```

3. Enable PFC on priorities 0 and 1.

```
(config-if)#priority-flow-control enable priority 0 1
```

4. Enable automatic priority flow control deadlock recovery mode timer with custom detection and recovery time parameters.

```
(config-if)#priority-flow-control deadlock recovery-mode timer detection-multiplier 10  
time-granularity 10 recovery-time 1000
```

PFC state XON mode

1. Set the IP address.

```
(config-if)#ip address 2.2.2.1/24
```

2. Enable the PFC. Configure the advertise flag and start sending DCBX TLVs in LLDP messages.

```
(config-if)#priority-flow-control mode on
```

3. Enable PFC on priorities 0 and 1.

```
(config-if)#priority-flow-control enable priority 0 1
```

4. Enable automatic priority flow control deadlock recovery mode timer with custom detection and recovery time parameters.

```
(config-if)#priority-flow-control deadlock recovery-mode timer detection-multiplier 10
time-granularity 10 recovery-time 1000
```

Global Mode

When any interface enters deadlock recovery mode, instead of allowing the deadlocked traffic to pass, traffic will be dropped if this command is set

```
(config)#priority-flow-control deadlock recovery-action drop
```

Validation

1. Verifying deadlock config and status for all interfaces.

```
#show priority-flow-control deadlock-status
```

```
Deadlock Detection and Recovery Configuration
```

```
-----
interface      recovery      detection      detection      recovery
                mode          multiplier     granularity    time
=====
eth1           Timer         10             10             1500
-----
```

```
Deadlock Detection and Recovery Status
```

```
-----
interface      pri    state          detection      last detection      last recovery
                state                                count             timestamp           timestamp
=====
eth1           1      deadlock       39            2025-05-29 19:03:49.481  -
-----
```

2. Verifying deadlock config and status for a specific interface

```
#show priority-flow-control deadlock-status interface eth1
```

```
Deadlock Detection and Recovery Configuration
```

```
-----
interface      recovery      detection      detection      recovery
                mode          multiplier     granularity    time
=====
eth1           Timer         10             10             1500
-----
```

Deadlock Detection and Recovery Status

interface	pri	state	detection count	last detection timestamp	last recovery timestamp
eth1	0	no deadlock	0	-	-
eth1	1	deadlock	35	2025-05-29 19:03:34.611	-
eth1	2	no deadlock	0	-	-
eth1	3	no deadlock	0	-	-
eth1	4	no deadlock	0	-	-
eth1	5	no deadlock	0	-	-
eth1	6	no deadlock	0	-	-
eth1	7	no deadlock	0	-	-

3. Clearing deadlock status for a specific interface

```
#clear priority-flow-control deadlock-status eth1
```

4. Clearing deadlock status for all interfaces.

```
#clear priority-flow-control deadlock-status
```

Using manual recovery on an interface

Once a deadlock is detected and no manual recovery mode is configured in the interface, it is possible to recover from the deadlock by manually entering and exiting recovery mode on supported boards with the below commands:

1. Start manual deadlock recovery on interface eth1.

```
#priority-flow-control eth1 deadlock manual-recovery start
```

2. Stop manual deadlock recovery on interface eth1.

```
#priority-flow-control eth1 deadlock manual-recovery stop
```

Validation**1. Verifying deadlock config and status for all interfaces.**

```
#show priority-flow-control deadlock-status
```

Deadlock Detection and Recovery Configuration

interface	recovery mode	detection multiplier	detection granularity	recovery time
eth1	Timer	10	10	1500

Deadlock Detection and Recovery Status

interface	pri	state	detection count	last detection timestamp	last recovery timestamp
eth1	1	deadlock	39	2025-05-29 19:03:49.481	-

2. Verifying deadlock config and status for a specific interface

```
#show priority-flow-control deadlock-status interface eth1
```

Deadlock Detection and Recovery Configuration

interface	recovery mode	detection multiplier	detection granularity	recovery time
eth1	Timer	10	10	1500

Deadlock Detection and Recovery Status

interface	pri	state	detection count	last detection timestamp	last recovery timestamp
eth1	0	no deadlock	0	-	-
eth1	1	deadlock	35	2025-05-29 19:03:34.611	-
eth1	2	no deadlock	0	-	-
eth1	3	no deadlock	0	-	-
eth1	4	no deadlock	0	-	-
eth1	5	no deadlock	0	-	-
eth1	6	no deadlock	0	-	-
eth1	7	no deadlock	0	-	-

3. Clearing deadlock status for a specific interface

```
#clear priority-flow-control deadlock-status eth1
```

4. Clearing deadlock status for all interfaces.

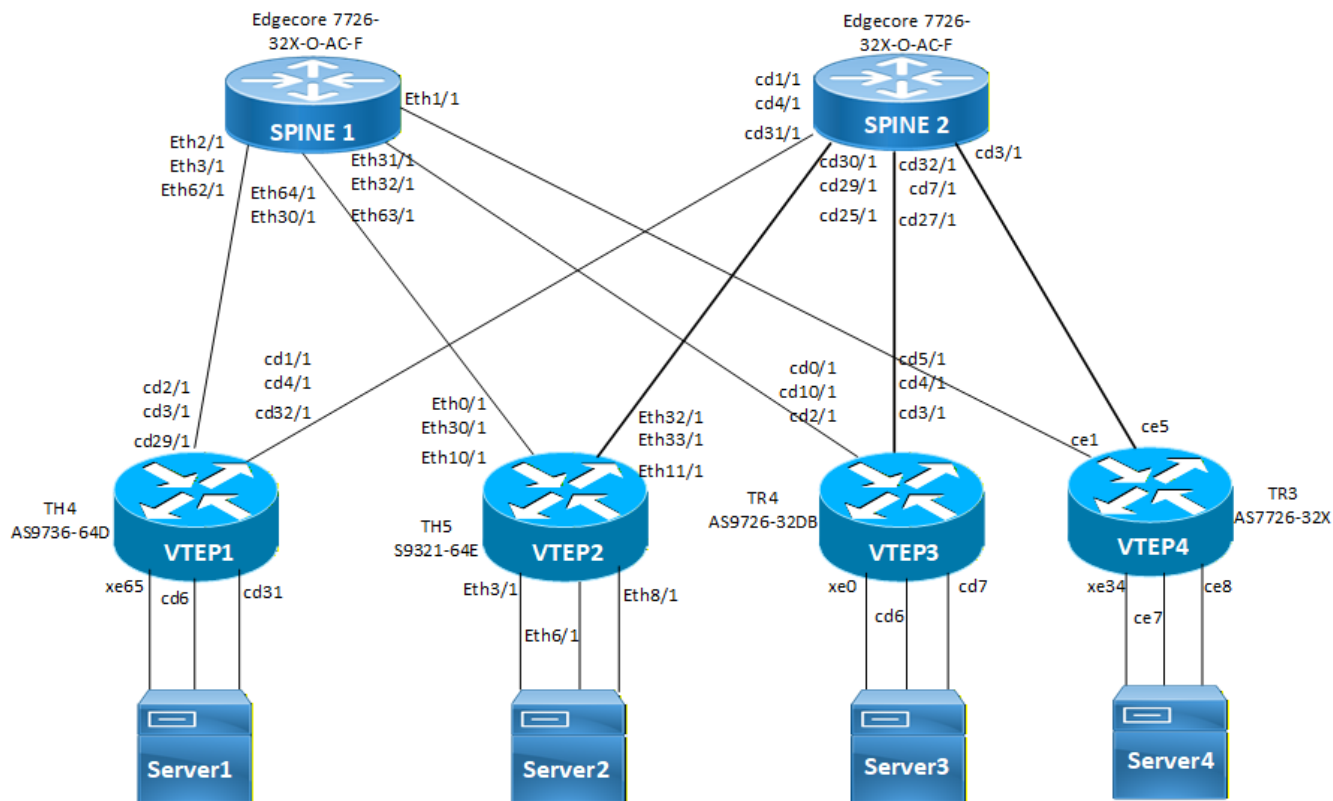
```
#clear priority-flow-control deadlock-status
```

Spine and Leaf Topology

The topology uses an EVPN-VXLAN Underlay using OSPF for the internal gateway protocol (IGP) and eBGP with unnumbered interfaces for advertising host and loopback routes. It also utilizes Dynamic Load Balancing (DLB) and Priority Flow Control (PFC) for PFC Dead Lock Detection requirements.

Leaf switches connect to both the spines layer (uplinks) and the end-hosts (downlinks, often GPUs/Servers).

Figure 4. PFC Dead Lock Detection and Recovery



Configure Leaf 1, Leaf 2, Leaf 3, Leaf 4, Spine 1 and Spine 2

The following configurations, enables PFC deadlock recovery on every priority interfaces on spines and leaves.

1. Configure global settings - PFC, QoS, DCB, DLB, and load balance rtag7 hashing.

```
data-center-bridging enable bridge 1priority-flow-control enable bridge 1
dynamic-load-balance enable
maximum-paths 64
load-balance rtag7 ipv4 dest-ipv4 src-ipv4 destl4-port srcl4-port protocol-id rocev2-dest-qpairs
load-balance rtag7 ipv6 dest-ipv6 src-ipv6 destl4-port srcl4-port rocev2-dest-qpairs next-hdr
policy-map type queuing default default-out-policy
class type queuing default q6
priority
lossless
exit
class type queuing default q7
priority
lossless
exit
```

2. Configure PFC default policy for all 8 traffic classes.

```
(config)#policy-map type queuing default default-out-policy
(config-cmap-que)#class type queuing default q0
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q1
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q2
```

```
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q3
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q4
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q5
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q6
(config-cmap-que)# priority
(config-cmap-que)# lossless
(config-cmap-que)# exit
(config)#class type queuing default q7
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
```

3. Configure PFC default lossless ECN policy.

```
(config)#policy-map type queuing default lossless_ecn_egress
(config-pmap-que-def)#class type queuing default q0
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
(config-pmap-que-def)#class type queuing default q1
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
(config-pmap-que-def)#class type queuing default q2
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
random-detect green min-threshold 500 max-threshold 600 yellow min-threshold 300 max-threshold 400
red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
(config-pmap-que-def)#class type queuing default q3
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
(config-pmap-que-def)#class type queuing default q4
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
(config-pmap-que-def)#class type queuing default q5
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
```

```
(config-pmap-que-def)#class type queuing default q6
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
(config-pmap-que-def)#class type queuing default q7
(config-pmap-c-que-def)#shape 10 gbps
(config-pmap-c-que-def)#priority
(config-pmap-c-que-def)#lossless
(config-pmap-c-que-def)#random-detect green min-threshold 500 max-threshold 600 yellow min-threshold
300 max-threshold 400 red min-threshold 100 max-threshold 200 packets ecn
(config-pmap-c-que-def)#exit
```

4. Configure routing map information to allow redistribution of routes. It is applied to incoming BGP routes, which sets the metric to 55555.

```
(config)#route-map HIG_MED permit 10
(config-route-map)#set metric 55555
```

5. Set the PFC mode to auto on the uplink interfaces.

On Leaf 1

```
(config)#interface cd1/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 10.1.12.1/24
(config-if)#ipv6 address 1012::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd3/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 10.1.11.1/24
(config-if)#ipv6 address 1011::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd2/1
(config-if)# priority-flow-control mode auto
(config-if)#ip address 20.1.11.1/24
(config-if)#ipv6 address 2011::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd4/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 20.1.12.1/24
(config-if)#ipv6 address 2012::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)# exit
```

```
(config)#interface cd29/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 30.1.11.1/24
(config-if)#ipv6 address 3011::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd32/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 30.1.12.1/24
(config-if)#ipv6 address 3012::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

On Leaf 2

```
(config)#interface ethernet0/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 10.1.13.1/24
(config-if)#ipv6 address 1013::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet32/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 10.1.14.1/24
(config-if)#ipv6 address 1014::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet30/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 20.1.13.1/24
(config-if)#ipv6 address 2013::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet33/1
(config-if)# priority-flow-control mode auto
(config-if)#ip address 20.1.14.1/24
(config-if)#ipv6 address 2014::1/64
(config-if)# shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)# ipv6 router ospf area 0.0.0.0 instance-id 0
```

```
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet10/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 30.1.13.1/24
(config-if)#ipv6 address 3013::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet11/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 30.1.14.1/24
(config-if)#ipv6 address 3014::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

On Leaf 3

```
(config)#interface cd2/1
(config-if)# priority-flow-control mode auto
(config-if)#ip address 10.1.16.1/24
(config-if)#ipv6 address 1016::1/64
(config-if)# shape rate 5 gbps burst 10000
(config-if)# ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)# ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)# lldp-agent
(config-if)# exit

(config)#interface cd0/1
(config-if)# priority-flow-control mode auto
(config-if)#ip address 10.1.15.1/24
(config-if)#ipv6 address 1015::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd10/1
(config-if)#priority-flow-control mode auto
(config-if)# ip address 20.1.15.1/24
(config-if)#ipv6 address 2015::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)# ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)# lldp-agent
(config-if)# exit

(config)#interface cd5/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 20.1.16.1/24
(config-if)#ipv6 address 2016::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
```

```
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd3/1
(config-if)# priority-flow-control mode auto
(config-if)#ip address 30.1.15.1/24
(config-if)#ipv6 address 3015::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd4/1
(config-if)#priority-flow-control mode auto
(config-if)#ip address 30.1.16.1/24
(config-if)#ipv6 address 3016::1/64
(config-if)#shape rate 5 gbps burst 10000
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

On Leaf 4

```
(config)#interface cel
(config-if)# priority-flow-control mode auto
(config-if)#ip address 10.1.17.1/24
(config-if)#ipv6 address 1017::1/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ce5
(config-if)# priority-flow-control mode auto
(config-if)#ip address 10.1.18.1/24
(config-if)# ipv6 address 1018::1/64
(config-if)# ip ospf cost 100
(config-if)# ipv6 ospf cost 100 instance-id 0
(config-if)# ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)# lldp-agent
(config-if)# exit
```

6. Set PFC mode and DeadLock Detection Recovery timer on downlink interfaces.

On Leaf 1

```
(config)#interface cd6
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 111.1.1.1/24
(config-if)#ipv6 address 111::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd31
(config-if)#description ***GPU ***
(config-if)#switchport
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 111-115
```

```
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#lldp-agent
(config-if)#exit

(config)#interface xe65
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 101.1.1.1/24
(config-if)#ipv6 address 101::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```



Note: The 'deadlock recovery-mode timer' command configured only on xe65 and not explicitly on cd6 or cd31.

On Leaf 2

```
(config)#interface ethernet3/1
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 102.1.1.1/24
(config-if)#ipv6 address 102::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet6/1
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 112.1.1.1/24
(config-if)#ipv6 address 112::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet8/1
(config-if)# description ***GPU ***
(config-if)# switchport
(config-if)# bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 211-215
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)# lldp-agent
(config-if)#exit
```

On Leaf 3

```
(config)#interface cd7
(config-if)# description ***GPU ***
(config-if)#switchport
```

```
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 311-315
(config-if)# priority-flow-control mode on
(config-if)# priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#lldp-agent
(config-if) exit

(config)#interface cd6
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 113.1.1.1/24
(config-if)#ipv6 address 113::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config-if)#interface xe0
(config-if)# description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)# priority-flow-control deadlock recovery-mode timer
(config-if)# priority-flow-control advertise-local-config
(config-if)# priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 103.1.1.1/24
(config-if)#ipv6 address 103::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)# lldp-agent
(config-if)# exit
```

On Leaf 4

```
(config)#interface ce7
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)# priority-flow-control advertise-local-config
(config-if)# priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 114.1.1.1/24
(config-if)# ipv6 address 114::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)# exit

(config)#interface ce8
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 124.1.1.1/24
(config-if)#ipv6 address 124::1/64
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface xe34
(config-if)#description ***GPU ***
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 104.1.1.1/24
(config-if)#ipv6 address 104::1/64
```

```
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

7. Configure interfaces on spines that connects to leaves.

On Spine 1

```
(config-if)#interface ethernet1/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 10.1.17.2/24
(config-if)#ipv6 address 1017::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config-if)#interface ethernet2/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 20.1.11.2/24
(config-if)#ipv6 address 2011::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet3/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 10.1.11.2/24
(config-if)#ipv6 address 1011::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config-if)#interface ethernet30/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 20.1.13.2/24
(config-if)#ipv6 address 2013::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config-if)#interface ethernet31/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
```

```
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 10.1.15.2/24
(config-if)#ipv6 address 1015::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet32/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 20.1.15.2/24
(config-if)#ipv6 address 2015::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet62/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 30.1.11.2/24
(config-if)#ipv6 address 3011::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface ethernet63/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 30.1.15.2/24
(config-if)#ipv6 address 3015::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
(config)#interface ethernet64/1
(config-if)# priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)# priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#load-interval 30
(config-if)#ip address 10.1.13.2/24
(config-if)#ipv6 address 1013::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)# exit
```

On Spine 2

```
(config)#interface cdl/1
```

```
(config-if)# priority-flow-control mode on
(config-if)# priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 10.1.12.2/24
(config-if)#ipv6 address 1012::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

```
(config)#interface cd3/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 10.1.18.2/24
(config-if)#ipv6 address 1018::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

```
(config)#interface cd4/1
(config-if)# priority-flow-control mode on
(config-if)# priority-flow-control deadlock recovery-mode timer
(config-if)# priority-flow-control advertise-local-config
(config-if)# priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 20.1.12.2/24
(config-if)# ipv6 address 2012::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)# ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)# lldp-agent
(config-if)# exit
```

```
(config)#interface cd7/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 20.1.16.2/24
(config-if)#ipv6 address 2016::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

```
(config)#interface cd25/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 30.1.14.2/24
(config-if)#ipv6 address 3014::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit
```

```
(config)#interface cd27/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
```

```

(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 30.1.16.2/24
(config-if)#ipv6 address 3016::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd29/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 20.1.14.2/24
(config-if)#ipv6 address 2014::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd30/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 10.1.14.2/24
(config-if)#ipv6 address 1014::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd31/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 30.1.12.2/24
(config-if)#ipv6 address 3012::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

(config)#interface cd32/1
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#priority-flow-control advertise-local-config
(config-if)#priority-flow-control enable priority 0 1 2 3 4 5 6 7
(config-if)#ip address 10.1.16.2/24
(config-if)#ipv6 address 1016::2/64
(config-if)#ip ospf cost 100
(config-if)#ipv6 ospf cost 100 instance-id 0
(config-if)#ipv6 router ospf area 0.0.0.0 instance-id 0
(config-if)#lldp-agent
(config-if)#exit

```

8. Configure the Underlay routing OSPF for Spine and Leaf interconnections.

On Leaf 1

```

(config)#router ospf 100
(config-router)#ospf router-id 1.1.1.1

```

```
(config-router)#network 1.1.1.1/32 area 0.0.0.0
(config-router)#network 10.1.11.0/24 area 0.0.0.0
(config-router)#network 10.1.12.0/24 area 0.0.0.0
(config-router)#network 20.1.11.0/24 area 0.0.0.0
(config-router)#network 20.1.12.0/24 area 0.0.0.0
(config-router)#network 30.1.11.0/24 area 0.0.0.0
(config-router)#network 30.1.12.0/24 area 0.0.0.0
(config-router)#network 101.1.1.0/24 area 0.0.0.0
(config-router)#network 111.1.1.0/24 area 0.0.0.0
(config-router)#network 171.1.1.0/24 area 0.0.0.0
(config-router)#network 171.1.2.0/24 area 0.0.0.0
(config-router)#network 171.1.3.0/24 area 0.0.0.0
(config-router)#network 171.1.4.0/24 area 0.0.0.0
(config-router)#network 171.1.5.0/24 area 0.0.0.0

(config)#router ipv6 ospf
(config-router)#router-id 1.1.1.1
```

On Leaf 2

```
(config)#router ospf 100
(config-router)#ospf router-id 2.2.2.2
(config-router)#network 2.2.2.2/32 area 0.0.0.0
(config-router)#network 10.1.13.0/24 area 0.0.0.0
(config-router)#network 10.1.14.0/24 area 0.0.0.0
(config-router)#network 20.1.13.0/24 area 0.0.0.0
(config-router)#network 20.1.14.0/24 area 0.0.0.0
(config-router)#network 30.1.13.0/24 area 0.0.0.0
(config-router)#network 30.1.14.0/24 area 0.0.0.0
(config-router)#network 40.1.13.0/24 area 0.0.0.0
(config-router)#network 50.1.13.0/24 area 0.0.0.0
(config-router)#network 60.1.13.0/24 area 0.0.0.0
(config-router)#network 102.1.1.0/24 area 0.0.0.0
(config-router)#network 112.1.1.0/24 area 0.0.0.0
(config-router)#network 172.1.1.0/24 area 0.0.0.0
(config-router)#network 172.1.2.0/24 area 0.0.0.0
(config-router)#network 172.1.3.0/24 area 0.0.0.0
(config-router)#network 172.1.4.0/24 area 0.0.0.0
(config-router)#network 172.1.5.0/24 area 0.0.0.0
```

On Leaf 3

```
(config)#router ospf 100
(config-router)#ospf router-id 3.3.3.3
(config-router)#network 3.3.3.3/32 area 0.0.0.0
(config-router)#network 10.1.15.0/24 area 0.0.0.0
(config-router)#network 10.1.16.0/24 area 0.0.0.0
(config-router)#network 20.1.15.0/24 area 0.0.0.0
(config-router)#network 20.1.16.0/24 area 0.0.0.0
(config-router)#network 30.1.15.0/24 area 0.0.0.0
(config-router)#network 30.1.16.0/24 area 0.0.0.0
(config-router)#network 103.1.1.0/24 area 0.0.0.0
(config-router)#network 113.1.1.0/24 area 0.0.0.0
(config-router)#network 173.1.1.0/24 area 0.0.0.0
(config-router)#network 173.1.2.0/24 area 0.0.0.0
(config-router)#network 173.1.3.0/24 area 0.0.0.0
(config-router)#network 173.1.4.0/24 area 0.0.0.0
(config-router)#network 173.1.5.0/24 area 0.0.0.0
```

On Leaf 4

```
(config)#router ospf 100
(config-router)#ospf router-id 4.4.4.4
(config-router)#network 4.4.4.4/32 area 0.0.0.0
(config-router)#network 10.1.17.0/24 area 0.0.0.0
(config-router)#network 10.1.18.0/24 area 0.0.0.0
(config-router)#network 104.1.1.0/24 area 0.0.0.0
(config-router)#network 114.1.1.0/24 area 0.0.0.0
```

```
(config-router)#network 174.1.1.0/24 area 0.0.0.0
(config-router)#network 174.1.2.0/24 area 0.0.0.0
(config-router)#network 174.1.3.0/24 area 0.0.0.0
(config-router)#network 174.1.4.0/24 area 0.0.0.0
(config-router)#network 174.1.5.0/24 area 0.0.0.0
```

On Spine 1

```
(config-router)#router ospf 100
(config-router)#ospf router-id 7.7.7.7
(config-router)#network 7.7.7.7/32 area 0.0.0.0
(config-router)#network 10.1.11.0/24 area 0.0.0.0
(config-router)#network 10.1.13.0/24 area 0.0.0.0
(config-router)#network 10.1.15.0/24 area 0.0.0.0
(config-router)#network 10.1.17.0/24 area 0.0.0.0
(config-router)#network 10.1.19.0/24 area 0.0.0.0
(config-router)#network 20.1.11.0/24 area 0.0.0.0
(config-router)#network 20.1.13.0/24 area 0.0.0.0
(config-router)#network 20.1.15.0/24 area 0.0.0.0
(config-router)#network 20.1.17.0/24 area 0.0.0.0
(config-router)#network 30.1.11.0/24 area 0.0.0.0
(config-router)#network 30.1.13.0/24 area 0.0.0.0
(config-router)#network 30.1.15.0/24 area 0.0.0.0
(config-router)#network 30.1.17.0/24 area 0.0.0.0
(config-router)#network 40.1.13.0/24 area 0.0.0.0
(config-router)#network 50.1.13.0/24 area 0.0.0.0
(config-router)#network 60.1.13.0/24 area 0.0.0.0
```

On Spine2

```
(config-router)#router ospf 100
(config-router)#ospf router-id 8.8.8.8
(config-router)#network 8.8.8.8/32 area 0.0.0.0
(config-router)#network 10.1.12.0/24 area 0.0.0.0
(config-router)#network 10.1.14.0/24 area 0.0.0.0
(config-router)#network 10.1.16.0/24 area 0.0.0.0
(config-router)#network 10.1.18.0/24 area 0.0.0.0
(config-router)#network 10.1.20.0/24 area 0.0.0.0
(config-router)#network 20.1.12.0/24 area 0.0.0.0
(config-router)#network 20.1.14.0/24 area 0.0.0.0
(config-router)#network 20.1.16.0/24 area 0.0.0.0
(config-router)#network 20.1.18.0/24 area 0.0.0.0
(config-router)#network 30.1.12.0/24 area 0.0.0.0
(config-router)#network 30.1.14.0/24 area 0.0.0.0
(config-router)#network 30.1.16.0/24 area 0.0.0.0
(config-router)#network 30.1.18.0/24 area 0.0.0.0
```

9. Configure the eBGP router in unnumbered mode on interfaces connecting to Spine and Leaf for advertising host and loopback.

On Leaf 1

```
(config)#router bgp 4294967201
(config-router)#bgp router-id 1.1.1.1
(config-router)#bgp log-neighbor-changes
(config-router)#neighbor underlay peer-group
(config-router)#neighbor underlay remote-as 4294967209
(config-router)#neighbor underlay shutdown
(config-router)#neighbor underlay authentication-key 0xb59db09d828b2528
(config-router)#neighbor underlay as-origination-interval 1
(config-router)#neighbor underlay advertisement-interval 0
(config-router)#neighbor underlay fall-over bfd
(config-router)#exit

(config-router)#bgp unnumbered-mode
(config-router-annum)#neighbor cd4/1 peergroup underlay
(config-router-annum)#neighbor cd3/1 peergroup underlay
(config-router-annum)#neighbor cd2/1 peergroup underlay
```

```
(config-router-unnum)#neighbor cd1/1 peergroup underlay
(config-router-unnum)#neighbor cd32/1 peergroup underlay
(config-router-unnum)#neighbor cd29/1 peergroup underlay
(config-router-unnum)#exit-unnumbered-mode

(config-router)#address-family ipv4 unicast
(config-router-af)#max-paths ebgp 10
(config-router-af)#redistribute connected
(config-router-af)#neighbor underlay activate

(config-router)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor cd4/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd3/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd2/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd1/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd32/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd29/1 route-map HIG_MED in
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-af)#exit-address-family
(config-router)#exit
```

On Leaf 2

```
(config)#router bgp 4294967202
(config-router)#bgp router-id 2.2.2.2
(config-router)#bgp log-neighbor-changes
(config-router)#neighbor underlay peer-group
(config-router)#neighbor underlay remote-as 4294967209
(config-router)#neighbor underlay shutdown
(config-router)#neighbor underlay authentication-key 0xb59db09d828b2528
(config-router)#neighbor underlay as-origination-interval 1
(config-router)#neighbor underlay advertisement-interval 0
(config-router)#neighbor underlay fall-over bfd

(config)#bgp unnumbered-mode
(config-router-unnum)#neighbor ethernet10/1 peergroup underlay
(config-router-unnum)#neighbor ethernet0/1 peergroup underlay
(config-router-unnum)#neighbor ethernet30/1 peergroup underlay
(config-router-unnum)#neighbor ethernet33/1 peergroup underlay
(config-router-unnum)#neighbor ethernet32/1 peergroup underlay
(config-router-unnum)#neighbor ethernet11/1 peergroup underlay
(config-router-unnum)#exit-unnumbered-mode

(config)#address-family ipv4 unicast
(config-router-af)#max-paths ebgp 10
(config-router-af)#redistribute connected
(config-router-af)#neighbor underlay activate

(config)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor ethernet10/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor ethernet0/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor ethernet30/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor ethernet33/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor ethernet32/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor ethernet11/1 route-map HIG_MED in
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-v4-unnum)#exit-address-family
(config-router-v4-af)#exit
```

On Leaf 3

```
(config)#router bgp 4294967203
(config-router)#bgp router-id 3.3.3.3
(config-router)#bgp log-neighbor-changes
(config-router)#neighbor underlay peer-group
(config-router)#neighbor underlay remote-as 4294967209
(config-router)#neighbor underlay shutdown
(config-router)#neighbor underlay authentication-key 0xb59db09d828b2528
```

```
(config-router)#neighbor underlay as-origination-interval 1
(config-router)#neighbor underlay advertisement-interval 0
(config-router)#neighbor underlay fall-over bfd

(config-router)#bgp unnumbered-mode
(config-router-unnum)#neighbor cd10/1 peergroup underlay
(config-router-unnum)#neighbor cd5/1 peergroup underlay
(config-router-unnum)#neighbor cd4/1 peergroup underlay
(config-router-unnum)#neighbor cd3/1 peergroup underlay
(config-router-unnum)#neighbor cd2/1 peergroup underlay
(config-router-unnum)#neighbor cd0/1 peergroup underlay
(config-router-unnum)#exit-unnumbered-mode

(config)#address-family ipv4 unicast
(config-router-af)#max-paths ebgp 10
(config-router-af)#redistribute connected
(config-router-af)#neighbor underlay activate

(config-router)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor cd10/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd5/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd4/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd3/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd2/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd0/1 route-map HIG_MED in
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-v4-unnum)#exit-address-family
(config-router-v4-af)#exit
```

On Leaf 4

```
(config)#router bgp 4294967204
(config-router)#bgp router-id 4.4.4.4
(config-router)#bgp log-neighbor-changes
(config-router)#neighbor underlay peer-group
(config-router)#neighbor underlay remote-as 4294967209
(config-router)#neighbor underlay shutdown
(config-router)#neighbor underlay authentication-key 0xb59db09d828b2528
(config-router)#neighbor underlay as-origination-interval 1
(config-router)#neighbor underlay advertisement-interval 0
(config-router)#neighbor underlay fall-over bfd

(config)#bgp unnumbered-mode
(config-router-unnum)#neighbor ce1 peergroup underlay
(config-router-unnum)# neighbor ce5 peergroup underlay
(config-router-unnum)# exit-unnumbered-mode

(config-router)#address-family ipv4 unicast
(config-router-af)#max-paths ebgp 10
(config-router-af)#redistribute connected
(config-router-af)#neighbor underlay activate

(config-router)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor ce1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor ce5 route-map HIG_MED in
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router-v4-af)#exit-address-family
```

On Spine1

```
(config)#router bgp 4294967209
(config-router)#bgp router-id 7.7.7.7
(config-router)#bgp log-neighbor-changes
(config-router)#neighbor underlay peer-group
(config-router)#neighbor underlay remote-as 1
(config-router)#neighbor underlay shutdown
(config-router)#neighbor underlay authentication-key 0xb59db09d828b2528
(config-router)#neighbor underlay as-origination-interval 1
```

```
(config-router)#neighbor underlay advertisement-interval 0
(config-router)#neighbor underlay fall-over bfd

(config-router)#bgp unnumbered-mode
(config-router-annum)#neighbor ethernet1/1 remote-as external
(config-router-annum)#neighbor ethernet1/1 peergroup underlay
(config-router-annum)#neighbor ethernet31/1 remote-as external
(config-router-annum)#neighbor ethernet31/1 peergroup underlay
(config-router-annum)#neighbor ethernet63/1 remote-as external
(config-router-annum)#neighbor ethernet63/1 peergroup underlay
(config-router-annum)#neighbor ethernet32/1 remote-as external
(config-router-annum)#neighbor ethernet32/1 peergroup underlay
(config-router-annum)#neighbor ethernet64/1 remote-as external
(config-router-annum)#neighbor ethernet64/1 peergroup underlay
(config-router-annum)#neighbor ethernet61/1 remote-as external
(config-router-annum)#neighbor ethernet61/1 peergroup underlay
(config-router-annum)#neighbor ethernet30/1 remote-as external
(config-router-annum)#neighbor ethernet30/1 peergroup underlay
(config-router-annum)#neighbor ethernet2/1 remote-as external
(config-router-annum)#neighbor ethernet2/1 peergroup underlay
(config-router-annum)#neighbor ethernet3/1 remote-as external
(config-router-annum)#neighbor ethernet3/1 peergroup underlay
(config-router-annum)#neighbor ethernet62/1 remote-as external
(config-router-annum)#neighbor ethernet62/1 peergroup underlay
(config-router-annum)#neighbor ethernet11/1 remote-as external
(config-router-annum)#neighbor ethernet11/1 peergroup underlay
(config-router-annum)#exit-unnumbered-mode

(config-router)#address-family ipv4 unicast
(config-router-af)#max-paths ebgp 10
(config-router-af)#redistribute connected
(config-router-af)#neighbor underlay activate

(config-router)#bgp v4-unnumbered-mode
(config-router-v4-annum)#neighbor ethernet1/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet31/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet63/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet32/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet64/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet61/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet30/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet2/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet3/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet62/1 route-map HIG_MED in
(config-router-v4-annum)#neighbor ethernet11/1 route-map HIG_MED in
(config-router-v4-annum)#exit-v4-unnumbered-mode
(config-router-v4-af)#exit-address-family
(config-router)#exit
```

On Spine2

```
(config)#router bgp 4294967209
(config-router)#bgp router-id 8.8.8.8
(config-router)#bgp log-neighbor-changes
(config-router)#neighbor underlay peer-group
(config-router)#neighbor underlay remote-as 1
(config-router)#neighbor underlay shutdown
(config-router)#neighbor underlay authentication-key 0xb59db09d828b2528
(config-router)#neighbor underlay as-origination-interval 1
(config-router)#neighbor underlay advertisement-interval 0
(config-router)#neighbor underlay fall-over bfd

(config-router)#bgp unnumbered-mode
(config-router-annum)#neighbor cd29/1 remote-as external
(config-router-annum)#neighbor cd29/1 peergroup underlay
(config-router-annum)#neighbor cd3/1 remote-as external
(config-router-annum)#neighbor cd3/1 peergroup underlay
(config-router-annum)#neighbor cd30/1 remote-as external
```

```
(config-router-unnum)#neighbor cd30/1 peergroup underlay
(config-router-unnum)#neighbor cd25/1 remote-as external
(config-router-unnum)#neighbor cd25/1 peergroup underlay
(config-router-unnum)#neighbor cd4/1 remote-as external
(config-router-unnum)#neighbor cd4/1 peergroup underlay
(config-router-unnum)#neighbor cd1/1 remote-as external
(config-router-unnum)#neighbor cd1/1 peergroup underlay
(config-router-unnum)#neighbor cd31/1 remote-as external
(config-router-unnum)#neighbor cd31/1 peergroup underlay
(config-router-unnum)#neighbor cd32/1 remote-as external
(config-router-unnum)#neighbor cd32/1 peergroup underlay
(config-router-unnum)#neighbor cd27/1 remote-as external
(config-router-unnum)#neighbor cd27/1 peergroup underlay
(config-router-unnum)#neighbor cd7/1 remote-as external
(config-router-unnum)#neighbor cd7/1 peergroup underlay
(config-router-unnum)#neighbor cd9/1 remote-as external
(config-router-unnum)#neighbor cd9/1 peergroup underlay
(config-router-unnum)#exit-unnumbered-mode

(config-router)#address-family ipv4 unicast
(config-router-af)# max-paths ebgp 10
(config-router-af)#redistribute connected
(config-router-af)#neighbor underlay activate

(config-router)#bgp v4-unnumbered-mode
(config-router-v4-unnum)#neighbor cd29/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd3/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd30/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd25/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd4/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd1/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd31/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd32/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd27/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd7/1 route-map HIG_MED in
(config-router-v4-unnum)#neighbor cd9/1 route-map HIG_MED in
(config-router-v4-unnum)#exit-v4-unnumbered-mode
(config-router)#exit-address-family
(config-router)#exit
```

Validation

Check the PFC details on Spine1 interfaces.

```
Spine1-TH5-7001#show priority-flow-control details all

Admin Configuration
-----
interface          mode  advertise willing  cap  link delay  priorities
                               allowance
=====
ethernet1/1        on   on       off    8    0           0 1 2 3 4 5 6 7
ethernet2/1        on   on       off    8    0           0 1 2 3 4 5 6 7
ethernet30/1       on   on       off    8    0           0 1 2 3 4 5 6 7
ethernet32/1       on   on       off    8    0           0 1 2 3 4 5 6 7
ethernet64/1       on   on       off    8    0           0 1 2 3 4 5 6 7
-----

Operational Configuration
-----
interface          state cap  link delay  priorities
                               allowance
=====
ethernet1/1        on   8    0           0 1 2 3 4 5 6 7
ethernet2/1        on   8    0           0 1 2 3 4 5 6 7
ethernet30/1       on   8    0           0 1 2 3 4 5 6 7
```

```

ethernet32/1      on      8      0      0 1 2 3 4 5 6 7
ethernet64/1     on      8      0      0 1 2 3 4 5 6 7

```

Change recovery-time to 100 and check the PFC deadlock status.

```
Spine1-TH5-7001#show priority-flow-control deadlock-status interface ethernet1/1
```

```
Deadlock Detection and Recovery Configuration
```

```

-----
interface          recovery      detection      detection      recovery
                   mode          multiplier     granularity    time
=====
ethernet1/1        XON           100            10              100
-----

```

```
Deadlock Detection and Recovery Status
```

```

-----
interface          pri    state          detection      last detection      last recovery
                   count          count          timestamp         timestamp
=====
ethernet1/1        0     no deadlock     0              -                   -
ethernet1/1        1     no deadlock     0              -                   -
ethernet1/1        2     no deadlock     0              -                   -
ethernet1/1        3     no deadlock     0              -                   -
ethernet1/1        4     no deadlock     0              -                   -
ethernet1/1        5     no deadlock     0              -                   -
ethernet1/1        6     no deadlock     0              -                   -
ethernet1/1        7     no deadlock     0              -                   -
-----

```

Remove recovery-time and check the PFC deadlock status.

```

#sh run int ethernet1/1
!
interface ethernet1/1
  description Connected-cel-7042
  priority-flow-control mode on
  priority-flow-control advertise-local-config
  priority-flow-control deadlock recovery-mode pfc-state-xon detection-multiplier 100 time-granularity
  10
  priority-flow-control enable priority 0 1 2 3 4 5 6 7
  load-interval 30
  ip address 104.1.1.2/24
  ipv6 address 1401::2/64
  mtu 9216
  ipv6 router ospf area 0.0.0.0 instance-id 0
!

```

No PFC deadlock detected.

```
Spine1-TH5-7001#show priority-flow-control deadlock-status
```

```
Deadlock Detection and Recovery Configuration
```

```

-----
interface          recovery      detection      detection      recovery
                   mode          multiplier     granularity    time
=====
ethernet1/1        XON           100            10              -
-----

```

```
Deadlock Detection and Recovery Status
```

```
-----  
----  
interface          pri    state      detection    last detection    last recovery  
                  count      count      count      timestamp        timestamp  
-----  
=====
```

PFC DD Commands

The following commands are introduced as part of the PFC DD recovery.

- [clear priority-flow-control deadlock-status \(page 156\)](#)
- [priority-flow-control deadlock manual-recovery \(page 157\)](#)
- [priority-flow-control deadlock recovery-action drop](#)
- [priority-flow-control deadlock recovery-mode timer \(page 158\)](#)
- [priority-flow-control deadlock recovery-mode timer \(page 158\)](#)
- [show priority-flow-control deadlock-status \(page 159\)](#)

clear priority-flow-control deadlock-status

Use this command to clear the PFC deadlock details for a specified interface or for all interfaces

Command Syntax

```
clear priority-flow-control deadlock-status [ IFNAME ]
```

Parameters

IFNAME

Name of the input or output interface

Default

None

Command Mode

Execution mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
#clear priority-flow-control deadlock-status interface eth1
```

priority-flow-control deadlock manual-recovery

Use this command to start/stop manually the PFC deadlock recovery on the specified interface.

Command Syntax

```
priority-flow-control <NAME> deadlock manual-recovery ( start | stop )
```

Parameters

IFNAME

Name of the input or output interface

Default

None

Command Mode

Execution mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
#priority-flow-control eth1 deadlock manual-recovery start  
#priority-flow-control eth1 deadlock manual-recovery stop
```

priority-flow-control deadlock recovery-action drop

Use this command to globally drop deadlocked traffic on Priority-based Flow Control (PFC) deadlock recovery. Use the no form of this command to allow deadlocked traffic when a PFC deadlock recovery occurs.

Command Syntax

```
priority-flow-control deadlock recovery-action drop
no priority-flow-control deadlock recovery-action drop
```

Parameters

None

Default

By default, PFC deadlocked traffic during a recovery is allowed.

Command Mode

Configure mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

#configure terminal (config)

```
#configure terminal (config)
#priority-flow-control deadlock recovery-action drop
```

priority-flow-control deadlock recovery-mode timer

Use this command to enable Priority-based Flow Control (PFC) deadlock and recovery on all priorities of an interface, using a timer to end the recovery phase.

Use the no form of this command to disable PFC deadlock detection and recovery on an interface.

Command Syntax

```
priority-flow-control deadlock recovery-mode timer [ detection-multiplier <1-1599> time-granularity  
<1|10|100> ] [ recovery-time <100-1599> ]  
no priority-flow-control deadlock recovery-mode
```

Parameters

detection-multiplier

Specify the detection multiplier duration in micro seconds.

time-granularit

Specify the time granularity duration in micro seconds.

recovery-time

Specify the Recovery time duration in micro seconds.

Default

By default, detection multiplier is 10, time granularity is 10ms and recovery time is 100ms.

PFC deadlock detection is disabled by default.

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
#configure terminal (config)  
(config)#interface xel  
(config-if)#priority-flow-control deadlock recovery-mode timer detection-multiplier 100  
time-granularity 100 recovery-time 1000
```

priority-flow-control deadlock recovery-mode pfc-state-xon

Use this command to enable Priority-based Flow Control (PFC) deadlock and recovery on all priorities of an interface, using XON packet reception end the recovery phase.

Use the `no` form of this command to disable PFC deadlock detection and recovery on an interface.

Command Syntax

```
priority-flow-control deadlock recovery-mode pfc-state-xon [ detection-multiplier <1-1599> time-  
granularity <1|10|100> ]  
no priority-flow-control deadlock recovery-mode
```

Parameters

detection-multiplier

Specify the detection multiplier duration in micro seconds.

time-granularity

Specify the time granularity duration in micro seconds.

Default

By default, detection multiplier is 10, time granularity is 10ms and recovery time is 100ms.

PFC deadlock detection is disabled by default.

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
#configure terminal (config)  
(config)#interface xel  
(config-if)#priority-flow-control deadlock recovery-mode  
pfc-state-xon detection-multiplier 100 time-granularity 100
```

show priority-flow-control deadlock-status

Use this command to display the PFC deadlock details for a specified interface or for all interfaces

Command Syntax

```
show priority-flow-control deadlock-status [ IFNAME ]
```

Parameters

IFNAME

Name of the input or output interface

Default

None

Command Mode

Execution mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
#show priority-flow-control deadlock-status

Deadlock Detection and Recovery Configuration
-----
interface          recovery    detection    detection    recovery
                   mode        multiplier   granularity   time
=====
xe0                 Timer      10           10           1500
-----

Deadlock Detection and Recovery Status
-----
----
interface          pri    state          detection    last detection    last recovery
                   pri    state          count        timestamp         timestamp
=====
xe0                 1     deadlock      39          2025-05-29 19:03:49.481  -
```

Implementation Examples

Use case for PFC monitoring:

In a cloud data center, RoCEv2 traffic (RDMA over Converged Ethernet) runs across the fabric. Lossless transmission is critical, and PFC is used to pause specific priorities when buffers approach congestion. Use PFC monitoring to detect:

- If too many pause frames are being sent (could indicate congestion hotspots).
- If pause frames are stuck (deadlock scenarios).

Use Case for ECN monitoring in Leaf-Spine Fabric:

A hyperscale data center enables ECN marking on switches to signal congestion without dropping packets. End-host TCP stacks respond by reducing transmission rates. For ECN monitoring:

- Enable ECN on switch interfaces.
- Monitor ECN-marked packets per flow.

Glossary

Key Terms/Acronym	Description
PFC	Priority-based Flow Control. A mechanism supported by OcNOS to pause frames using defined times for each of the eight priority classes to prevent congestion.
XOFF	A control signal sent from the receiver to the transmitter, indicating that the receiver is congested and cannot accept additional data. It is signaled by a non-zero pause time in the PFC frame.
XON	A control signal sent from the receiver to the transmitter, indicating readiness to accept data (a no-pause condition).
Timer Mode	An automatic recovery mode where the system clears the deadlock after a user-defined time interval (recovery-time). This is the only mode supported by Trident3 (TR3) platforms.
PFC-State-XON Mode	An automatic recovery mode where recovery ends when the interface receives a PFC XON frame, signaling the pause condition is lifted.
Manual Mode	A recovery option that requires explicit user action via CLI commands to start and stop the recovery phase.

PFC Deadlock Detection and Recovery

Overview

Priority-based Flow Control (PFC) helps manage traffic in networks by pausing specific flows during congestion. However, under certain conditions, a deadlock can occur when cyclical dependencies between flows create a loop of PFC pause events that prevents traffic from making forward progress indefinitely.

Priority Flow Control Pause Frames

PFC uses the standard pause frame mechanism with an additional 14 bytes of padding in the frame. This padding contains a 2-byte value for each of the eight priority classes, specifying the pause time in quanta for that class.

Example for priority pause frame:

```
Pause Frame:
Control Opcode: 0x0101 (Priority Pause)
Pause Time (8 priority classes):
Class 0: 0x0000 (no pause)
```

```
Class 1: 0x1234 (pause time in quanta)
Class 2: 0x5678 (pause time in quanta)
...
Class 7: 0x9ABC (pause time in quanta)
```

In the example above, the **pause time in quanta** field defines if the pause frame has XON or XOFF set for that class:

- **XON (X-On):** A control signal sent from the receiver to the transmitter to indicate readiness to accept data. For example, Class 0 represents a no-pause condition.
- **XOFF (X-Off):** A control signal sent from the receiver to the transmitter indicating that it cannot accept additional data due to congestion. For example, Classes 1, 2, and 7 above specify a non-zero pause time (in quanta), signaling the transmitter to temporarily halt transmission.

Workflow of PFC Frames

- **Transmission:** The transmitter sends data to the receiver.
- **XOFF:** The receiver sends an XOFF signal to the transmitter, indicating that it is congested and cannot process more data.
- **Pause:** The transmitter pauses sending data to the receiver.
- **XON:** The receiver sends an XON signal to the transmitter, indicating that it is ready to receive data again.
- **Resume:** The transmitter resumes sending data to the receiver.

PFC deadlock:

A deadlock may occur when the receiver continuously sends XOFF signals for one or more classes, preventing the transmitter from sending any traffic. This feature is designed to detect such deadlocks and initiate recovery mechanisms.

To handle such critical situation, the OcNOS system provides PFC Deadlock Detection and Recovery capability. This chapter describes how to:

- Enable PFC deadlock detection and recovery on a specific interface
 - [Configure Using Timer mode \(page 149\)](#)
 - [PFC State XON mode \(page 152\)](#)
- Configure the global PFC deadlock detection and recovery action to drop
 - [Global Mode \(page 153\)](#)

Feature Characteristics

Deadlock Detection

- The system monitors PFC queues for extended periods in the XOFF state.
- If a queue remains paused beyond a configurable threshold, a deadlock event is declared.
- An interrupt is raised to inform software of the detected deadlock.

Deadlock Recovery

Once a deadlock is detected, software moves the affected queue into an ignore PFC XOFF state, allowing traffic scheduling to resume.

Recovery can be configured on a per-interface basis and supports three modes:

- **Timer Mode:** Recovery ends automatically after a user-defined time interval. The system then clears the interrupt and restarts the detection timer. It is an automatic recovery method and recovery starts after a configurable detection-multiplier times time-granularity period. During that period, traffic will be allowed by default, but can also be dropped if the configuration priority-flow-control deadlock recovery-action drop is set. Recovery also ends automatically after a optionally configurable recovery-time period.



Note: Traffic will gradually decrease to zero if the `recovery-mode timer` is not configured; otherwise, it will continue indefinitely.

- **PFC-State-XON Mode:** Recovery ends when the interface receives a PFC XON frame, signaling that the pause condition is lifted.
- **Manual Mode:** Recovery requires explicit user action with CLI commands. This option is only valid if no automatic recovery mode is configured.

Limitation:

- Manual recovery mode is not supported in Trident3 (TD3) platforms or Tomahawk3 (TH3) platforms.
- Trident3 (TD3) platforms support deadlock recovery only in timer mode.
- Trident3 (TD3) platforms do not support 1ms time granularity.
- Tomahawk2 (TH2) platforms are not supported.

Benefits

- Prevents indefinite traffic stalls due to PFC loops.
- Provides flexible recovery options (automatic or manual).
- Improves network reliability in environments that rely on PFC.

Prerequisites

- The device should be enabled with PFC.

Configuring PFC Deadlock Detection and Recovery

Topology

This topology illustrates Deadlock condition using single Switch1:

1. the `xe0` interfaces receives the regular traffic flow from a traffic source which is PFC enabled.
2. the same priority 1 regular unicast traffic is transmitted to the `xe1` interface which is also PFC enabled.
3. the traffic source sends pause frames that have XOFF on priority 1 to the `xe0` interface.
4. the regular traffic received on `xe0` interface is slowed down until no more packets arrive from the regular traffic flow.

Figure 5. PFC Deadlock Detection and Recovery

Configure Using Timer mode

The following configurations, enables PFC deadlock recovery on every priority interfaces. Optionally, the detection multiplier, time granularity and recovery time parameters may be set, otherwise the defaults (respectively 10, 10, 100) will be used.

Execute the following steps to configure PFC on both interfaces.

1. Create a bridge 1 as an IEEE VLAN-enabled bridge. Enable DCB and PFC on the bridge.

```
(config)#bridge 1 protocol ieee vlan-bridge
(config)#data-center-bridging enable bridge 1
(config)#priority-flow-control enable bridge 1
```

2. Enable PFC priorities 0 and 1 on both the interfaces **xe0** and **xe1**.

```
(config)#interface xe0
(config-if)#switchport
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan all
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control enable priority 0 1
(config-if)#load-interval 30

(config)interface xe1
(config-if)#switchport
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan all
(config-if)#priority-flow-control mode on
(config-if)#priority-flow-control enable priority 0 1
(config-if)#load-interval 30
```

3. Configure the priority flow control

deadlock recovery mode timer on **xe0** interface.

```
(config)#interface xe0
(config-if)#priority-flow-control deadlock recovery-mode timer
(config-if)#commit
(config-if)#
```

4. Configure the PFC deadlock detection happens with non-default detection-multiplier and time-granularity parameters

```
(config)#interface xe0
(config-if)#priority-flow-control deadlock recovery-mode timer detection-multiplier 10 time-granularity 10 recovery-time 1000
```

Validation

Check the PFC priority 1 traffic received and send on **xe0** and **xe1** interfaces.

```
OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
| xe0       | 0.00    | 0      | 1176.62 | 147078 |
| xe1       | 1176.62 | 147077 | 0.00    | 0      |
OcNOS#
```

Check the queue-level statistics on each interface to monitor transmitted traffic, drops, and queue.

```
OcNOS#clear interface counters
OcNOS#show interface counters queue-stats
```

E - Egress, I - Ingress, Q-Size is in bytes
 * indicates monitor is active

```

+-----+-----+-----+-----+-----+-----+-----+
| Interface | Queue/Class-map | Q-Size | Tx pkts | Tx bytes | Dropped pkts | Dropped bytes |
+-----+-----+-----+-----+-----+-----+-----+
xe0 q1 (E) 0 427127 427127000 0 0
    
```

After 30 seconds, check link throughput, should be like this:

```

OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0         0.00      0      1176.62  147077
xe1        1176.62   147077    0.00     0
OcNOS#
    
```

Check the traffic gradually stops after starting the XOFF P1 traffic to simulate the deadlock.

```

OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0         0.66     162     540.53  67566
xe1        1176.62   147077    0.11    211
    
```

Check the TX traffic on **xe0** is completely stopped after 30 seconds.

```

OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0         1.23     300     0.00     0
xe1        1176.62   147077    0.20    397
    
```

Check the traffic flow is restored after 30 seconds even though the XOFF packets exists.

```

OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0         1.23     300     1176.68  147086
xe1        1176.59   147074    0.00     0
    
```

Check the PFC deadlock is detected:

```

OcNOS#show priority-flow-control deadlock-status interface xe0

Deadlock Detection and Recovery Configuration
-----
interface      recovery      detection      detection      recovery
              mode         multiplier    granularity    time
=====
xe0             Timer         10            10             1500
-----

Deadlock Detection and Recovery Status
-----
interface      pri    state      detection      last detection      last recovery
              state  count      count          timestamp           timestamp
=====
xe0            0     no deadlock  0              -                   -
xe0            1     deadlock   35            2025-05-29 19:03:34.611  -
    
```

```
xe0          2    no deadlock          0    -    -
xe0          3    no deadlock          0    -    -
xe0          4    no deadlock          0    -    -
xe0          5    no deadlock          0    -    -
xe0          6    no deadlock          0    -    -
xe0          7    no deadlock          0    -    -
```

OcNOS#show priority-flow-control deadlock-status

Deadlock Detection and Recovery Configuration

```
-----
interface      recovery      detection      detection      recovery
                mode          multiplier     granularity    time
=====
xe0             Timer         10             10             1500
-----
```

Deadlock Detection and Recovery Status

```
-----
interface      pri    state          detection      last detection      last recovery
                state          count           timestamp          timestamp
=====
xe0             1      deadlock       39             2025-05-29 19:03:49.481
=====
```

Check the deadlock detection count:

OcNOS#show priority-flow-control deadlock-status interface xe0

Deadlock Detection and Recovery Configuration

```
-----
interface      recovery      detection      detection      recovery
                mode          multiplier     granularity    time
=====
xe0             Timer         1              1              1500
-----
```

Deadlock Detection and Recovery Status

```
-----
interface      pri    state          detection      last detection      last recovery
                state          count           timestamp          timestamp
=====
xe0             0      no deadlock       0    -    -
xe0             1      no deadlock      163   2025-05-29 19:55:16.842  2025-05-29
19:55:18.344
xe0             2      no deadlock       0    -    -
xe0             3      no deadlock       0    -    -
xe0             4      no deadlock       0    -    -
xe0             5      no deadlock       0    -    -
xe0             6      no deadlock       0    -    -
xe0             7      no deadlock       0    -    -
=====
```

Change recovery-time to 1000 and check the PFC deadlock.

OcNOS#show priority-flow-control deadlock-status interface xe0

Deadlock Detection and Recovery Configuration

```
-----
interface      recovery      detection      detection      recovery
                mode          multiplier     granularity    time
=====
```

```
xe0          Timer          10          10          1000
```


 Deadlock Detection and Recovery Status

interface	pri	state	detection count	last detection timestamp	last recovery timestamp
xe0	0	no deadlock	0	-	-
xe0	1	deadlock	9	2025-05-29 20:04:23.251	-
xe0	2	no deadlock	0	-	-
xe0	3	no deadlock	0	-	-
xe0	4	no deadlock	0	-	-
xe0	5	no deadlock	0	-	-
xe0	6	no deadlock	0	-	-
xe0	7	no deadlock	0	-	-

PFC State XON mode

Execute the following steps to configure PFC on both interfaces.

1. Create a bridge 1 as an IEEE VLAN-enabled bridge. Enable DCB on the bridge.

```
(config)#bridge 1 protocol ieee vlan-bridge
(config)#data-center-bridging enable bridge 1
```

2. Enables PFC on the bridge. Configure the advertise flag and start sending DCBX TLVs in LLDP messages.

```
(config)#priority-flow-control enable bridge 1
(config-if)#priority-flow-control mode on
```

3. Enable PFC on priorities 0 and 1.

```
(config-if)#priority-flow-control enable priority 0 1
```

4. Enable automatic priority flow control deadlock recovery mode `pfc-state-xon` with custom detection time.

```
(config-if)#priority-flow-control deadlock recovery-mode pfc-state-xon detection-multiplier 10 time-granularity 10
```

Validation

Change the recovery-mode to `pfc-state-xon` and check PFC deadlocks

```
OcNOS#show int counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
xe0          1.23      300      1176.62  147078
xe1          1176.62   147077   0.00     0
```

```
OcNOS##show priority-flow-control deadlock-status
```

Deadlock Detection and Recovery Configuration

interface	recovery mode	detection multiplier	detection granularity	recovery time
xe0	XON	10	10	-

 Deadlock Detection and Recovery Status

interface	pri	state	detection count	last detection timestamp	last recovery timestamp
xe0	1	deadlock	1	2025-05-29 20:12:58.089	-

Global Mode

When any interface enters deadlock recovery mode, instead of allowing the deadlocked traffic to pass, traffic will be dropped if this command is set globally.

```
(config)#priority-flow-control deadlock recovery-action drop
```

Validation

Change the global action command to drop and check the effect on traffic during the recovery process:

Interface	Rx mbps	Rx pps	Tx mbps	Tx pps
xe0	1.23	300	0.00	0
xe1	1176.62	147077	0.00	0

OcNOS#

Manual Recovery

Once a deadlock is detected and no manual recovery mode is configured in the interface, it is possible to recover from the deadlock by manually entering and exiting recovery mode on supported boards with the below commands:

1. Start manual deadlock recovery on interface eth1.

```
#priority-flow-control xe0 deadlock manual-recovery start
```

2. Stop manual deadlock recovery on interface eth1

```
#priority-flow-control xe1 deadlock manual-recovery stop
```

Validation

Start sending XOFF on P1 manual recovery mode and check if traffic starts to flow:

```
OcNOS#priority-flow-control xe0 deadlock-manual-recovery start
OcNOS#show int count rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
xe0         | 1.21    | 295    | 311.12  | 38890   |
xe1         | 1176.60 | 147075 | 0.15    | 287     |
OcNOS#show int count rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
xe0         | 1.21    | 295    | 311.12  | 38890   |
xe1         | 1176.60 | 147075 | 0.15    | 287     |
OcNOS#show int count rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
xe0         | 1.21    | 296    | 595.44  | 74430   |
xe1         | 1176.60 | 147074 | 0.10    | 191     |
```

Stop manual recovery mode and check if traffic stops.

```
OcNOS#priority-flow-control xe0 deadlock-manual-recovery stop
OcNOS#show int count rate mbps
```

```

+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0          1.22      297      823.27   102909
xe1          1176.61    147076   0.06     115
OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0          1.22      297      823.27   102909
xe1          1176.61    147076   0.06     115
OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0          1.21      296      710.29   88786
xe1          1176.61    147075   0.08     147
OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0          1.21      296      710.29   88786
xe1          1176.61    147075   0.08     147
OcNOS#show int count rate mbps
+-----+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+-----+
xe0          1.22      297      537.98   67248
xe1          1176.61    147076   0.11     208
OcNOS#

```

Clear PFC Deadlock

Following are the two commands to clear the deadlock.

1. Clearing deadlock status for a specific interface

```
#clear priority-flow-control deadlock-status xe1
```

2. Clearing deadlock status for all interfaces

```
#clear priority-flow-control deadlock-status
```

Show Running Configuration

```

OcNOS#show running-config
!
! Software version: UFI_S9110-32X-OcNOS-DC-NA-7.0.0.999- 08/21/2025 14:08:11
!
! Last configuration change at 05:02:36 UTC Sat Aug 16 2025 by root
!
!
service password-encryption
!
logging console 0
logging monitor 7
logging level nsm 7
logging level hsl 7
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
qos enable
!
bridge 1 protocol rstp vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1

```

```
snmp-server enable snmp vrf management
snmp-server view all .1 included vrf management
snmp-server community test group network-operator vrf management
snmp-server host 192.168.5.18 traps version 2c test vrf management host-vrf management
feature dns relay
ip dns relay
ipv6 dns relay
!
ip access-list 1
 10 permit ahp any any
!
policy-map type queuing default default-out-policy
 class type queuing default q0
   priority
   lossless
   exit
 class type queuing default q1
   priority
   lossless
   exit
 class type queuing default q2
   priority
   lossless
   exit
 class type queuing default q3
   priority
   lossless
   exit
 class type queuing default q4
   priority
   lossless
   exit
 class type queuing default q5
   priority
   lossless
   exit
 class type queuing default q7
   priority
   lossless
   exit
!
vlan database
 vlan-reservation 4029-4094
 vlan 10 bridge 1
!
ip vrf management
!
interface ce0
!
interface ce6
!
interface ce7
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
 priority-flow-control mode on
 priority-flow-control enable priority 0 1
 load-interval 30
!
interface ce8
 switchport
 bridge-group 1
 switchport mode trunk
 switchport trunk allowed vlan all
 priority-flow-control mode on
 priority-flow-control enable priority 0 1
 load-interval 30
```

```
!  
interface ce9  
!  
interface ce31  
!  
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 xe32  
  switchport  
  bridge-group 1  
  switchport mode trunk  
  switchport trunk allowed vlan all  
  priority-flow-control mode on  
  priority-flow-control enable priority 0 1  
  load-interval 30  
!  
  exit  
!  
line console 0  
  exec-timeout 35791  
line vty 0 3  
  exec-timeout 35791  
!  
!  
end
```

CLI Commands

The feature introduces the following configuration commands.

clear priority-flow-control deadlock-status

Use this command to clear the PFC deadlock details for a specified interface or for all interfaces.

Command Syntax

```
clear priority-flow-control deadlock-status [ IFNAME ]
```

Parameters

IFNAME

Name of the input or output interface.

Default

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
#clear priority-flow-control deadlock-status interface eth1
```

priority-flow-control deadlock manual-recovery

Use this command to start/stop manually the PFC deadlock recovery on the specified interface.

Command Syntax

```
priority-flow-control <NAME> deadlock manual-recovery ( start | stop )
```

Parameters

IFNAME

Name of the input or output interface.

Default

Nones

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
#priority-flow-control eth1 deadlock manual-recovery start  
#priority-flow-control eth1 deadlock manual-recovery stop
```

priority-flow-control deadlock recovery-action drop

Use this command to globally drop deadlocked traffic on Priority-based Flow Control (PFC) deadlock recovery.

Use the `no` form of this command to allow deadlocked traffic when a PFC deadlock recovery occurs.

Command Syntax

```
priority-flow-control deadlock recovery-action drop  
no priority-flow-control deadlock recovery-action drop
```

Parameters

None

Default

By default, PFC deadlocked traffic during a recovery is allowed.

Command Mode

Configuration mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
#configure terminal (config)
#priority-flow-control deadlock recovery-action drop
```

priority-flow-control deadlock recovery-mode timer

Use this command to enable Priority-based Flow Control (PFC) deadlock and recovery on all priorities of an interface, using a timer to end the recovery phase.

Use the `no` form of this command to disable PFC deadlock detection and recovery on an interface.

Command Syntax

```
priority-flow-control deadlock recovery-mode timer
[ detection-multiplier <1-1599> time-granularity <1|10|100> ]
[ recovery-time <100-1599> ]

no priority-flow-control deadlock recovery-mode
```

Parameters

detection-multiplier

Specify the detection multiplier duration in micro seconds.

time-granularity

Specify the time granularity duration in micro seconds.

recovery-time

Specify the Recovery time duration in micro seconds.

Default

By default, detection multiplier is 10, time granularity is 10ms and recovery time is 100ms.

PFC deadlock detection is disabled by default.

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
(config)#interface xe1
(config-if)#priority-flow-control deadlock recovery-mode
timer detection-multiplier 100 time-granularity 100
recovery-time 1000
```

```
OcNOS(config)#interface xe0
OcNOS(config-if)#no priority-flow-control deadlock recovery-mode
```

priority-flow-control deadlock recovery-mode pfc-state-xon

Use this command to enable Priority-based Flow Control (PFC) deadlock and recovery on all priorities of an interface, using XON packet reception end the recovery phase.m

Use the `no` form of this command to disable PFC deadlock detection and recovery on an interface.

Command Syntax

```
priority-flow-control deadlock recovery-mode pfc-state-xon
[ detection-multiplier <1-1599> time-granularity <1|10|100> ]
```

```
no priority-flow-control deadlock recovery-mode
```

Parameters

detection-multiplier

Specify the detection multiplier duration in micro seconds.

time-granularity

Specify the time granularity duration in micro seconds.

Default

By default, detection multiplier is 10, and time granularity is 10ms.

PFC deadlock detection is disabled by default.

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

#configure terminal (config)

```
#configure terminal (config)
(config)#interface xe1

(config-if)#priority-flow-control deadlock recovery-mode pfc-state-xon
detection-multiplier 100 time-granularity 100
```

show priority-flow-control deadlock-status

Use this command to display the PFC deadlock details for a specified interface or for all interfaces.

Command Syntax

```
show priority-flow-control deadlock-status [ IFNAME ]
```

Parameters**IFNAME**

Name of the input or output interface.

Default

None

Command Mode

Exec mode

Applicability

This command was introduced before OcNOS version 7.0.0.

Example

```
#show priority-flow-control deadlock-status

Deadlock Detection and Recovery Configuration
-----
interface recovery detection detection recovery
mode multiplier granularity time
=====
xe0 Timer 10 10 1500

-----

Deadlock Detection and Recovery Status
-----
-----
interface pri state detection last detection last recovery
count timestamp timestamp
=====
xe0 1 deadlock 39 2025-05-29 19:03:49.481 -
```

Implementation Examples**Use case for PFC monitoring:**

In a cloud data center, RoCEv2 traffic (RDMA over Converged Ethernet) runs across the fabric. Lossless transmission is critical, and PFC is used to pause specific priorities when buffers approach congestion. Use PFC monitoring to detect:

- If too many pause frames are being sent (could indicate congestion hotspots).
- If pause frames are stuck (deadlock scenarios).

Use Case for ECN monitoring in Leaf-Spine Fabric:

A hyperscale data center enables ECN marking on switches to signal congestion without dropping packets. End-host TCP stacks respond by reducing transmission rates. For ECN monitoring:

- Enable ECN on switch interfaces.
- Monitor ECN-marked packets per flow.

Glossary

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

Key Terms/Acronym	Description
PFC	Priority-based Flow Control. A mechanism to pause specific flows during congestion using pause frames based on defined times for each priority class.
XOFF	A control signal sent from the receiver to the transmitter, indicating that the receiver is congested and cannot accept additional data. It is signaled by a non-zero pause time in the PFC frame.
XON	A control signal sent from the receiver to the transmitter, indicating readiness to accept data (a no-pause condition).
Timer Mode	An automatic recovery mode where the system clears the deadlock after a user-defined time interval (recovery-time). This is the only mode supported by Trident3 (TR3) platforms.
PFC-State-XON Mode	An automatic recovery mode where recovery ends when the interface receives a PFC XON frame, signaling the pause condition is lifted.
Manual Mode	A recovery option that requires explicit user action via CLI commands to start and stop the recovery phase.

PFC Frames and ECN Packets Monitoring

Overview

OcNOS supports Priority-based Flow Control (PFC) to pause frames using defined times for each of the eight priority classes. This prevents congestion and improves transmission performance by letting the transmitter adjust its data flow according to the receiver's processing capacity.

Also supports Explicit Congestion Notification (ECN), which provides end-to-end congestion signaling between ECN-enabled senders and receivers in TCP/IP networks. Instead of dropping packets, ECN marks them to indicate congestion, prompting the sender to temporarily reduce its transmission rate until congestion clears. This reduces both packet loss and delay. ECN is defined in RFC 3168.

Feature Characteristics

This functionality enables:

- ECN marked packet monitoring on an interface
- PFC paused frames monitoring on an interface
- Monitored interfaces generate logs, NETCONF notifications, and SNMP traps whenever monitored packets are detected, including PFC frames and ECN-marked packets.

Limitation:

This functionality is applicable to the chips Tomahawk 2 (TH2) series platforms, Tomahawk3 (TH3) platforms, Tomahawk4 (TH4) platforms, Tomahawk5 (TH5) platforms, Trident3 (TR3) platforms and Trident4 (TR4) platforms.

Benefits

Improved Congestion Management – Prevents buffer overflows and packet drops by dynamically controlling traffic flow.

Per-Priority Traffic Control – Ensures that critical traffic classes (e.g., storage or real-time applications) are not impacted by congestion in other classes.

Reduced Packet Loss – Uses packet marking instead of dropping to signal congestion, minimizing retransmissions.

Higher Throughput Efficiency – Link utilization can be optimized via adjusting transmission rates based on real-time network conditions.

Prerequisites

PFC monitoring data requires a working PFC configuration and active PFC traffic. Similarly, ECN monitoring data requires a working ECN configuration and active ECN traffic.

Configuring PFC Frames and ECN Packets Monitoring

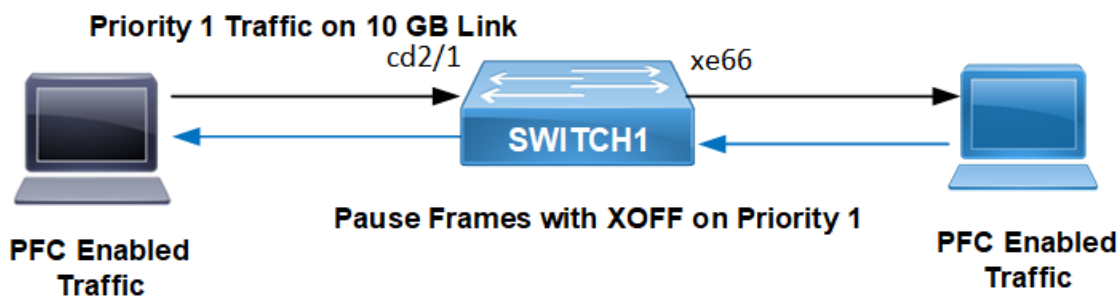
The configuration procedure outlines the steps required to enable ECN and PFC Support for Lossless TCP/IP Transport on the L3 networks, ensuring the network can handle high-priority, lossless AI/ML traffic.

Topology

The topology uses a Switch1 with an ingress interface *cd2/1* (connected to a node which generates traffic) and an egress interface *xe66* (connected to destination node which receives the traffic). Congestion is induced on the egress interface *xe66* using shapers within QoS policy maps.


The following topology shows PFC pause frame monitoring on the ingress and egress interfaces of Switch 1.

Figure 6. PFC Enabled Bridge



The following topology shows ECN Marked packets monitoring on the ingress and egress interfaces of Switch 1.

Figure 7. ECN Enabled Bridge



Configuration for ECN Marking and PFC Pausing

The following steps configure global setting for monitoring of PFC and ECN packets transmitted and received on an interface when monitoring is enabled in a Layer 3 routed scenario.



Note: Before configuration meet all [Prerequisites \(page 162\)](#).

1. Configure global settings - QoS, VLAN/Bridge, ingress port *cd2/1* as a trunk port in bridge 1, allowing VLAN 2 and egress port on Switch 1.

```
(config)#qos enable

(config)#vlan database
(config-vlan)#vlan-reservation 4001-4094
(config-vlan)#vlan 2 bridge 1 state enable

(config)#bridge 1 protocol rstp vlan-bridge
```

```
(config)#interface cd2/1
(config-if)#description Switch1
(config-if)#switchport
(config-if)#bridge-group 1
(config-if)#switchport mode trunk
(config-if)#switchport trunk allowed vlan add 2
(config-if)#load-interval 30
(config-if)#mtu 9216
```

ECN Configuration (Routing and ECN Marking)

Perform the following setup to enable ECN marking on congestion.

2. Enable ingress L3 Switch Virtual Interface (SVI) for Vlan1.2 with IP 10.1.1.1/24.

```
(config)#interface vlan1.2
(config-if)#ip address 10.1.1.1/24
(config-if)#mtu 9216
```

3. Configure egress port xe66 as a routed interface where congestion is applied.

```
(config)#interface xe66
(config-if)#description Connected-Destination Node
(config-if)#load-interval 30
(config-if)#ip address 20.1.1.1/24
(config-if)#mtu 9216
(config-if)#service-policy type queuing output ECN
(config-if)#monitor ecn
```

4. Configure OSPF router to ensure reachability between the networks connected to vlan1.2 and xe66.

```
(config)#router ospf 100
(config-router)#ospf router-id 1.1.1.1
(config-router)#network 1.1.1.1/32 area 0.0.0.0
(config-router)#network 10.1.1.0/24 area 0.0.0.0
(config-router)#network 20.1.1.0/24 area 0.0.0.0
```

5. Configure ECN policy map with queues *q0*, *q1*, *q2*, *q4*, *q5* for priority lossless and *q3* with shape 1 gbps to induce congestion, *priority*, and *random-detect, packets ecn* to enable ECN marking based on WRED thresholds.

```
(config)#policy-map type queuing default ECN
(config-cmap-que)#class type queuing default q0
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q1
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q2
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q3
(config-cmap-que)#shape 1 gbps
(config-cmap-que)#priority
(config-cmap-que)#random-detect green min-threshold 40 max-threshold 50 yellow min-threshold 70 max-
threshold 80 red min-threshold 100 max-threshold 110 packets ecn
(config-cmap-que)#exit
(config)#class type queuing default q4
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q5
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
```

6. Apply the ECN policy map to the egress interface xe66 using service-policy type queuing output ECN. Enable monitor ecn on xe66 to generate system logs for ECN marking events. Ensure PFC is not applied when only ECN is required.

```
(config)#interface xe66
(config-if)#description Egress-interface
(config-if)#load-interval 30
(config-if)#ip address 20.1.1.1/24
(config-if)#mtu 9216
(config-if)#service-policy type queuing output ECN
(config-if)#monitor ecn
```

PFC Configuration (Routing and PFC Pausing)

Perform the following setup to enable PFC pausing instead of ECN marking on congestion.

7. Configure PFC policy.

```
(config)#policy-map type queuing default PFC
(config-cmap-que)#class type queuing default q0
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q1
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q2
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q3
(config-cmap-que)#shape 1 gbps
(config-cmap-que)#priority
(config-cmap-que)#random-detect green min-threshold 40 max-threshold 50 yellow min-threshold 70 max-
threshold 80 red min-threshold 100 max-threshold 110
(config-cmap-que)#exit
(config)#class type queuing default q4
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
(config)#class type queuing default q5
(config-cmap-que)#priority
(config-cmap-que)#lossless
(config-cmap-que)#exit
```

8. Enable PFC policy.

```
OcNOS(config)#interface cd2/1
OcNOS(config-if)#switchport
OcNOS(config-if)#bridge-group 1
OcNOS(config-if)#switchport mode trunk
OcNOS(config-if)#switchport trunk allowed vlan all
OcNOS(config-if)#priority-flow-control mode on
OcNOS(config-if)#priority-flow-control enable priority 0 1 2 3 4
OcNOS(config-if)#load-interval 30
OcNOS(config-if)#mtu 9216
OcNOS(config-if)#monitor pfc
OcNOS(config-if)#shape rate 1 gbps burst 1000
```

9. Apply PFC policy.

```
OcNOS(config)#interface xe66
OcNOS(config-if)#description Egress-port
OcNOS(config-if)# priority-flow-control mode on
OcNOS(config-if)#priority-flow-control enable priority 0 1 2 3 4
OcNOS(config-if)#load-interval 30
OcNOS(config-if)#ip address 20.1.1.1/24
```

```
OcNOS(config-if)#mtu 9216
OcNOS(config-if)#service-policy type queuing output PFC
```

Sample Show Running Configuration on Switch

For ECN

```
!
service password-encryption
!
logging console 5
logging monitor 5
logging level all 7
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
qos enable
!
hostname DUT1
port cd2 breakout 4X10g
bridge 1 protocol rstp vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
feature dns relay
ip dns relay
ipv6 dns relay
!
policy-map type queuing default ECN
class type queuing default q0
priority
lossless
exit
class type queuing default q1
priority
lossless
exit
class type queuing default q2
priority
lossless
exit
class type queuing default q3
shape 1 gbps
priority
random-detect green min-threshold 40 max-threshold 50 yellow min-threshold 70 max-threshold 80 red
min-threshold 100 max-threshold 110 packets ecn
exit
class type queuing default q4
priority
lossless
exit
class type queuing default q5
priority
lossless
exit
!
vlan database
vlan-reservation 4001-4094
vlan 2 bridge 1 state enable
!
ip vrf management
!
interface cd1
!
```

```
interface cd2/1
  description Connected-STC
  switchport
  bridge-group 1
  switchport mode trunk
  switchport trunk allowed vlan add 2
  load-interval 30
  mtu 9216
!
interface cd64
!
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
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface vlan1.2
  ip address 10.1.1.1/24
  mtu 9216
!
interface xe65
!
interface xe66
  description Connected-DUT2
  load-interval 30
  ip address 20.1.1.1/24
  mtu 9216
  service-policy type queuing output ECN
  monitor ecn
!
  exit
!
router ospf 100
  ospf router-id 1.1.1.1
  network 1.1.1.1/32 area 0.0.0.0
  network 10.1.1.0/24 area 0.0.0.0
  network 20.1.1.0/24 area 0.0.0.0
!
router bgp 65001
  bgp router-id 1.1.1.1
  bgp log-neighbor-changes
  neighbor underlay peer-group
  neighbor underlay remote-as 65001
  neighbor underlay authentication-key 0x52211cdd013e0b79
  neighbor underlay as-origination-interval 1
  neighbor underlay advertisement-interval 0
  neighbor underlay fall-over bfd
!
  bgp unnumbered-mode
  neighbor xe66 peergroup underlay
  exit-unnumbered-mode
!
  address-family ipv4 unicast
  max-paths ebgp 10
  max-paths ibgp 64
  redistribute connected
  neighbor underlay activate
!
  bgp v4-unnumbered-mode
```

```
exit-v4-unnumbered-mode
!
exit-address-family
!1
exit
!
!
end

!
#
```

For PFC

```
!
service password-encryption
!
logging console 5
logging monitor 5
logging level all 7
!
!
snmp-server enable traps link linkDown
snmp-server enable traps link linkUp
!
qos enable
!
hostname DUT1
port cd2 breakout 4X10g
bridge 1 protocol rstp vlan-bridge
tfo Disable
errdisable cause stp-bpdu-guard
data-center-bridging enable bridge 1
feature dns relay
ip dns relay
ipv6 dns relay
!
policy-map type queuing default PFC
class type queuing default q0
priority
lossless
exit
class type queuing default q1
priority
lossless
exit
class type queuing default q2
priority
lossless
exit
class type queuing default q3
shape 1 gbps
priority
exit
class type queuing default q4
priority
lossless
exit
class type queuing default q5
priority
lossless
exit
!
vlan database
vlan-reservation 4001-4094
vlan 2 bridge 1 state enable
!
```

```
ip vrf management
!
interface cd1
!
interface cd2/1
  description Connected-STC
  switchport
  bridge-group 1
  switchport mode trunk
  switchport trunk allowed vlan add 2
  priority-flow-control mode on
  priority-flow-control enable priority 0 1 2 3 4
  load-interval 30
  mtu 9216
  monitor pfc
!
interface cd64
!
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
!
interface lo.management
  ip vrf forwarding management
  ip address 127.0.0.1/8
  ipv6 address ::1/128
!
interface vlan1.2
  ip address 10.1.1.1/24
  mtu 9216
!
interface xe65
!
interface xe66
  description Connected-DUT2
  priority-flow-control mode on
  priority-flow-control enable priority 0 1 2 3 4
  load-interval 30
  ip address 20.1.1.1/24
  mtu 9216
  service-policy type queuing output PFC
  monitor ecn
!
  exit
!
router ospf 100
  ospf router-id 1.1.1.1
  network 1.1.1.1/32 area 0.0.0.0
  network 10.1.1.0/24 area 0.0.0.0
  network 20.1.1.0/24 area 0.0.0.0
!
exit
!
!
end
!
```

Validation

ECN Validation

Verify the traffic rates on interfaces.

```
#sh int counters rate mbps
+-----+-----+-----+-----+-----+
|   Interface   |   Rx mbps  |   Rx pps  |   Tx mbps  |   Tx pps  |
+-----+-----+-----+-----+-----+
cd2/1           | 4325.03    | 4223658   | 0.00       | 1         |
xe66           | 0.00       | 5         | 1000.12    | 1008187  |
switch1#
switch1#sh int counters rate mbps
+-----+-----+-----+-----+-----+
|   Interface   |   Rx mbps  |   Rx pps  |   Tx mbps  |   Tx pps  |
+-----+-----+-----+-----+-----+
cd2/1           | 4325.03    | 4223658   | 0.00       | 1         |
xe66           | 0.00       | 5         | 1000.12    | 1008187  |
switch1#
```

Interface *cd2/1* (ingress L3 port leading to SVI *vlan1.2*) shows a high receive rate (~4325 Mbps), while the egress interface *xe66* shows a transmit rate capped at ~1000 Mbps (1 Gbps) due to the applied shaper. This confirms congestion is occurring on *xe66*.

Verify packet and byte counters for traffic passing through queues defined in applied policy maps.

```
#sh policy-map statistics
Type qos class-map statistics:
+-----+-----+-----+-----+-----+
|   Class-map   | Match pkts | Match bytes | Dropped |
| pkts | Dropped Bytes |
+-----+-----+-----+-----+-----+
Type queuing class-map statistics:
+-----+-----+-----+-----+-----+
|   Class-map   | Total pkts | Total bytes | Dropped pkts | Dropped
| Bytes |
+-----+-----+-----+-----+-----+
cd2/1
q6           | 213        | 33228      | 0          | 0
q7           | 1          | 68         | 0          | 0
xe66
q3           | 208204280  | 25817330720 | 661503879  | 84672496512
q7           | 23         | 2308       | 0          | 0
Switch1#
Switch1#sh policy-map statistics
Type qos class-map statistics:
+-----+-----+-----+-----+-----+
|   Class-map   | Match pkts | Match bytes | Dropped |
| pkts | Dropped Bytes |
+-----+-----+-----+-----+-----+
Type queuing class-map statistics:
+-----+-----+-----+-----+-----+
|   Class-map   | Total pkts | Total bytes | Dropped pkts | Dropped
| Bytes |
+-----+-----+-----+-----+-----+
-----+
```

```

cd2/1
q6          214          33384          0          0
q7          1           68            0          0
xe66
q3          209081096      25926055904    664719044    85084037632
q7          23           2308          0          0
Switch1 #

```

For interface `xe66`, queue `q3` shows a very large number of dropped packets (`664719044`) alongside significant total packets (`209081096`) and total bytes (`25926055904`) successfully transmitted. As ECN marking via *random-detect* is enabled on this queue, these dropped packets (`664719044`) actually represent ECN-marked packets, not physical discards. The thresholds are configured to mark packets rather than drop them upon congestion.

Verify the count of packets marked with ECN CE (Congestion Experienced) code point on a per-interface basis.

```

Switch1#sh int counters ecn
+-----+
| Interface          | ECN marked packets |
+-----+
xe66                  168391233

Switch1#sh int counters ecn
+-----+
| Interface          | ECN marked packets |
+-----+
xe66                  170407629

Switch1#sh int counters ecn
+-----+
| Interface          | ECN marked packets |
+-----+
xe66                  171415809

Switch1#2025 Oct 28 11:39:57.895 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface -
xe66, ECN MARKED PKT: 5040864
2025 Oct 28 11:40:02.895 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040864
2025 Oct 28 11:41:47.900 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040864
2025 Oct 28 11:41:52.901 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040936
2025 Oct 28 11:41:57.901 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040936
2025 Oct 28 11:42:02.901 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040864
2025 Oct 28 11:42:07.901 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040864
2025 Oct 28 11:42:12.901 : Switch1 : HSL : NOTIF : [IF_ECN_MONITOR_4]: ECN: Interface - xe66, ECN
MARKED PKT: 5040900

```

Interface `xe66` shows a large and steadily increasing number of ECN marked packets (`171415809`). This directly confirms that the ECN mechanism is actively marking packets on the congested egress interface as configured.

These log messages are generated due to the *monitor ecn* command on `xe66`. The logs periodically report the cumulative count of ECN Marked packet on interface `xe66`, providing real-time visibility into the ECN marking activity. The counts align with the increasing values seen in *sh int counters ecn*.

The configuration given in the [Validation \(page 170\)](#) successfully sets up an L3 path, induces congestion on the egress interface `xe66` via shaping, and applies an ECN policy. Validation confirms that packets exceeding the WRED thresholds in *queue 3* are being marked with ECN (not dropped), as shown by the dedicated ECN counters, interpreted policy map statistics, and system logs.

PFC Validation

Verify the traffic rates on interfaces.

```
Switch1 #sh int counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
cd2/1      | 1032.38 | 1008180 | 14.34   | 28006   |
xe66      | 0.00    | 5       | 1000.12 | 1008184 |
Switch1 #sh int counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
cd2/1      | 1032.38 | 1008180 | 14.34   | 28006   |
xe66      | 0.00    | 5       | 1000.12 | 1008184 |
Switch1 #sh int counters rate mbps
+-----+-----+-----+-----+
| Interface | Rx mbps | Rx pps | Tx mbps | Tx pps |
+-----+-----+-----+-----+
cd2/1      | 1032.38 | 1008179 | 14.34   | 28006   |
xe66      | 0.00    | 5       | 1000.12 | 1008187 |
```

The egress interface *xe66* is capped at ~ 1000.12 Mbps due to the shaper. The ingress interface *cd2/1* also shows a receive rate throttled down to ~ 1032.38 Mbps. This indicates that PFC pausing is being applied upstream from *cd2/1*.

Verify the counters for PFC pause frames sent and received per priority per interface.

```
Switch1 #sh priority-flow-control statistics all
interface      pri  pause sent  pause received
=====
cd2/1          0    0           0
cd2/1          1    0           0
cd2/1          2    0           0
cd2/1          3    3732510    0
cd2/1          4    0           0
cd2/1          5    0           0
cd2/1          6    0           0
cd2/1          7    0           0
xe66           0    0           0
xe66           1    0           0
xe66           2    0           0
xe66           3    0           0
xe66           4    0           0
xe66           5    0           0
xe66           6    0           0
xe66           7    0           0
Switch1 #
Switch1 #sh priority-flow-control statistics all
interface      pri  pause sent  pause received
=====
cd2/1          0    0           0
cd2/1          1    0           0
cd2/1          2    0           0
cd2/1          3    3760522    0
cd2/1          4    0           0
cd2/1          5    0           0
cd2/1          6    0           0
cd2/1          7    0           0
xe66           0    0           0
xe66           1    0           0
xe66           2    0           0
xe66           3    0           0
xe66           4    0           0
xe66           5    0           0
xe66           6    0           0
xe66           7    0           0
Switch1 #
Switch1 #sh priority-flow-control statistics all
interface      pri  pause sent  pause received
```

```

=====
cd2/1          0    0          0
cd2/1          1    0          0
cd2/1          2    0          0
cd2/1          3   3816534    0
cd2/1          4    0          0
cd2/1          5    0          0
cd2/1          6    0          0
cd2/1          7    0          0
xe66           0    0          0
xe66           1    0          0
xe66           2    0          0
xe66           3    0          0
xe66           4    0          0
xe66           5    0          0
xe66           6    0          0
xe66           7    0          0
Switch1 #

Switch1 #sh priority-flow-control statistics all
interface      pri   pause sent   pause received
=====
cd2/1          0    0          0
cd2/1          1    0          0
cd2/1          2    0          0
cd2/1          3   3844539    0
cd2/1          4    0          0
cd2/1          5    0          0
cd2/1          6    0          0
cd2/1          7    0          0
xe66           0    0          0
xe66           1    0          0
xe66           2    0          0
xe66           3    0          0
xe66           4    0          0
xe66           5    0          0
xe66           6    0          0
xe66           7    0          0
Switch1#

Switch1#sh priority-flow-control statistics all
interface      pri   pause sent   pause received
=====
cd2/1          0    0          0
cd2/1          1    0          0
cd2/1          2    0          0
cd2/1          3   3928546    0
cd2/1          4    0          0
cd2/1          5    0          0
cd2/1          6    0          0
cd2/1          7    0          0
xe66           0    0          0
xe66           1    0          0
xe66           2    0          0
xe66           3    0          0
xe66           4    0          0
xe66           5    0          0
xe66           6    0          0
xe66           7    0          0
Switch1 #

```

Interface *cd2/1* shows a large and rapidly increasing count of *pause sent* frames specifically for *priority 3* (~3732510 -> ~3928546). No pause frames are received (pause received is 0). This confirms that Switch1 is sending PFC pause frames out of the ingress interface *cd2/1* for **priority 3**. This happens because the downstream path (egress interface *xe66*) is congested for *queue 3* (due to the shaper), and PFC is enabled for this priority.

Verify the administrative and operational status of PFC per interface.

```
Switch1 #sh priority-flow-control deatails all
```

Admin Configuration

```
-----
interface          mode  advertise willing  cap  link delay  priorities
                    allowance
=====
cd2/1              on   on      off    8    0          0 1 2 3 4
xe66               on   on      off    8    0          0 1 2 3 4
-----
```

Operational Configuration

```
-----
interface          state cap  link delay  priorities
                    allowance
=====
cd2/1              on   8    0          0 1 2 3 4
xe66               on   8    0          0 1 2 3 4
-----
```

```
Switch1 #sh priority-flow-control details all
```

Admin Configuration

```
-----
interface          mode  advertise willing  cap  link delay  priorities
                    allowance
=====
cd2/1              on   on      off    8    0          0 1 2 3 4
xe66               on   on      off    8    0          0 1 2 3 4
-----
```

Operational Configuration

```
-----
interface          state cap  link delay  priorities
                    allowance
=====
cd2/1              on   8    0          0 1 2 3 4
xe66               on   8    0          0 1 2 3 4
-----
```

```
Switch1 #
```

```
Switch1 #sh priority-flow-control details all
```

Admin Configuration

```
-----
interface          mode  advertise willing  cap  link delay  priorities
                    allowance
=====
cd2/1              on   on      off    8    0          0 1 2 3 4
xe66               on   on      off    8    0          0 1 2 3 4
-----
```

Operational Configuration

```
-----
interface          state cap  link delay  priorities
                    allowance
=====
cd2/1              on   8    0          0 1 2 3 4
xe66               on   8    0          0 1 2 3 4
-----
```

```
Switch1 #
```

```
Switch1 #sh priority-flow-control details all
```

Admin Configuration

```
-----
interface          mode  advertise willing  cap  link delay  priorities
```

```

=====
                                     allowance
=====
cd2/1          on  on      off      8      0      0 1 2 3 4
xe66          on  on      off      8      0      0 1 2 3 4
=====

Operational Configuration
-----
interface      state cap  link delay  priorities
                                     allowance
=====
cd2/1          on   8     0           0 1 2 3 4
xe66          on   8     0           0 1 2 3 4
=====

Switch1 #

Switch1 #2025 Oct 28 11:48:37.913 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface -
cd2/1, PG[3]: Pause-Tx: 140027
2025 Oct 28 11:48:42.913 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140026
~2025 Oct 28 11:48:47.913 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140028
2025 Oct 28 11:48:52.913 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140023
2025 Oct 28 11:48:57.914 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140023
2025 Oct 28 11:49:02.914 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140025
2025 Oct 28 11:49:07.914 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140026
2025 Oct 28 11:49:12.914 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140023
2025 Oct 28 11:49:17.914 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140037
2025 Oct 28 11:49:22.915 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140014
2025 Oct 28 11:49:27.915 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140031
2025 Oct 28 11:49:32.915 : Switch1 : HSL : NOTIF : [IF_PFC_MONITOR_4]: PFC: Interface - cd2/1, PG
[3]: Pause-Tx: 140019

```

Both *cd2/1* and *xe66* shows Admin Configuration mode and Operational Configuration state as *on* for *priorities 0-4*. This confirms PFC is active and negotiated correctly on the relevant interfaces for the configured priority.

Logs periodically report the number of Pause-Tx (transmitted pause frames) for Priority Group 3 (PG[3]) on interface *cd2/1*, confirming the PFC activity shown in the statistics command. Conversely, if pause frames are received rather than transmitted, equivalent Pause-Rx logs will be displayed.

The configuration given in the [ECN Validation \(page 170\)](#) establishes an L3 path with shaping on egress (*xe66*) and PFC enabled on both ingress (*cd2/1*) and egress for relevant priorities. Validation confirms that congestion on the egress interface triggers PFC pause frames to be sent from the ingress interface (*cd2/1*), successfully throttling the traffic source and preventing packet loss due to the egress shaper.

PFC-ECN Commands

The following commands are introduced as part of the PFC and ECN monitoring.

- [monitor ecn \(page 178\)](#)
- [monitor pfc \(page 178\)](#)

monitor ecn

Use this command to enable Explicit-Congestion-Notification (ECN) marked packets monitoring on a physical interface.

Use the `no` form of this command to disable ECN monitoring on the interface

Command Syntax

```
monitor ecn
no monitor ecn
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
(config)#interface xe1
(config-if)#monitor ecn
```

monitor pfc

Use this command to enable Priority-based Flow Control (PFC) pause frames monitoring on a physical interface.

Use the `no` form of this command to disable PFC monitoring on the interface.

Command Syntax

```
monitor pfc
no monitor pfc
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
(config)#interface xel
(config-if)#monitor pfc
```

Glossary

Key Terms/Acronym	Description
PFC	Priority-based Flow Control. A mechanism supported by OcNOS to pause frames using defined times for each of the eight priority classes to prevent congestion.
ECN	Explicit Congestion Notification. A mechanism defined in RFC 3168 that provides end-to-end congestion signaling between ECN-enabled senders and receivers in TCP/IP networks. Instead of dropping packets, ECN marks them to indicate congestion.

PFC Frames and ECN Packets Monitoring**Overview**

OcNOS supports Protocols Supported for DCBX to pause frames using defined times for each of the eight priority classes. This prevents congestion and improves transmission performance by letting the transmitter adjust its data flow according to the receiver's processing capacity.

Also supports Explicit Congestion Notification (ECN), which provides end-to-end congestion signaling between ECN-enabled senders and receivers in TCP/IP networks. Instead of dropping packets, ECN marks them to indicate congestion, prompting the sender to temporarily reduce its transmission rate until congestion clears. This reduces both packet loss and delay. ECN is defined in RFC 3168.

Feature Characteristics

This functionality enables:

- ECN marked packet monitoring on an interface
- PFC paused frames monitoring on an interface
- Monitored interfaces generate logs, NETCONF notifications, and SNMP traps whenever monitored packets are detected, including PFC frames and ECN-marked packets.

Limitation:

This functionality is applicable to the chips Tomahawk2 (TH2) platforms , Tomahawk3 (TH3) platforms ,Tomahawk4 (TH4) platforms Tomahawk5 (TH5) platforms, , Trident3 (TD3) platforms and Trident4 (TD4) platforms.

Benefits

Improved Congestion Management – Prevents buffer overflows and packet drops by dynamically controlling traffic flow.

Per-Priority Traffic Control – Ensures that critical traffic classes (e.g., storage or real-time applications) are not impacted by congestion in other classes.

Reduced Packet Loss – Uses packet marking instead of dropping to signal congestion, minimizing retransmissions.

Higher Throughput Efficiency – Link utilization can be optimized via adjusting transmission rates based on real-time network conditions.

Prerequisites

PFC monitoring data requires a working PFC configuration and active PFC traffic. Similarly, ECN monitoring data requires a working ECN configuration and active ECN traffic.

CLI Commands

The feature introduces the following configuration commands.

[monitor ecn](#)

[monitor pfc](#)

monitor ecn

Use this command to enable Explicit-Congestion-Notification (ECN) marked packets monitoring on a physical interface.

Use the `no` form of this command to disable ECN monitoring on the interface

Command Syntax

```
monitor ecn
no monitor ecn
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
(config)#interface xel
(config-if)#monitor ecn
```

monitor pfc

Use this command to enable Priority-based Flow Control (PFC) pause frames monitoring on a physical interface.

Use the `no` form of this command to disable PFC monitoring on the interface.

Command Syntax

```
monitor pfc
no monitor pfc
```

Parameters

None

Default

None

Command Mode

Interface mode

Applicability

This command is introduced in OcNOS version 7.0.0.

Example

```
(config)#interface xe1
(config-if)#monitor pfc
```

Glossary

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

Key Terms/Acronym	Description
PFC	Priority-based Flow Control. A mechanism supported by OcNOS to pause frames using defined times for each of the eight priority classes to prevent congestion.
ECN	Explicit Congestion Notification. A mechanism defined in RFC 3168 that provides end-to-end congestion signaling between ECN-enabled senders and receivers in TCP/IP networks. Instead of dropping packets, ECN marks them to indicate congestion.

ECN and PFC Support for Lossless VxLAN Transport

OcNOS 7.0 enables Explicit Congestion Notification (ECN) and Priority Flow Control (PFC) operation over VxLAN overlays, allowing operators to extend lossless transport capabilities across multi-tenant AI fabrics and frontend network.

These enhancements provides:

- Scalable Layer 2 and Layer 3 multi-tenancy.
- End-to-end lossless transport across overlay networks.
- Seamless integration of AI workload isolation with high-performance GPU fabric requirements.

For more details, refer to the [Unified ECN and PFC Support for Lossless VxLAN Transport](#) section in the *OcNOS Virtual Extensible LAN Guide*, Release 7.0.0.

Switch Packet Buffer Tuning

This release introduces Network Switch Packet Buffer Tuning, a system designed to enhance network switch performance by avoiding congestion and packet drops. This feature allows for the allocation of packet buffer size based on traffic priority classes, known as Priority Groups (PGs), instead of physical ports.

Key Enhancements Include:

- Custom device responses to Priority-based Flow Control (PFC) pause storms, enabling precise control over when the switch transmits pause frames to prevent packet loss.
- Priority Group (PG) configuration with specific limits on shared memory and the ability to set PFC X-OFF and X-ON offsets to trigger pause frames during congestion.
- Queue-specific buffer limits using a dynamic threshold (alpha value) for fine-grained control over buffer consumption from the shared pool.
- Global adjustment of buffer limits, simplifying configuration.

Supported Platforms: This feature is intended for LTSW chipsets (Tomahawk4 (TH4) platforms, Tomahawk5 (TH5) platforms, Trident4 (TD4) platforms) and DC chipsets (Tomahawk3 (TH3) platforms, Trident3 (TD3) platforms).

For more details, refer to the [Switch Packet Buffer Tuning](#) section in *OcNOS Quality of Service Guide*, Release 7.0.0.

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.

Mirror Filtered Data to CPU	183
NetConf Access Control Model User Guide	184
Overview	184
Feature Characteristics	184
Configuration	185
Common NACM Rule Fields	186
Creating NetConf RPC for NACM	189
Glossary	198
sFlow - Sample Packet Monitoring for Multiple Interfaces	199
Overview	199
Features Characteristics	200
Benefits	200
Prerequisites	200
Configuration	200
Configuring sFlow with User Defined VRFs	203
Implementation Examples	206
Commands	212
Troubleshooting	215
VxLAN OAM for Overlay Networks	216
CLI-Script and CLI-Shell	217
Overview	217
Limitations	217
Configuration	218
Validation	219
CLI-Script and CLI-Shell Commands	221
System Limits and Counters	234
Overview	234
System Limits and Counters Limitation	235
System Limits and Counters Configuration	235
System Limits and Counters Implementation Example	242
System Limits and Counters Commands	242
System Limit Counters Troubleshooting	252
System Limit Counters Glossary	253

Mirror Filtered Data to CPU

Mirroring to CPU with filter feature provides the ability to mirror filtered data plane packets to CPU. It enables sniffing of selected packets that match the programmed filter condition and real-time monitoring in the Network Operating System.

This feature enables monitoring in the switching devices, such as leaf and spine switches. Monitoring at the leaf provides visibility into north-south traffic (between endpoints and external networks or services) and monitoring at the spine provides visibility into east-west traffic, i.e., between leaf switches.

For more details, refer to the *Mirror Filtered Data to CPU* topic in the OcNOS Layer 2 Guide, Release 7.0.0.

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>
  <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> <ul style="list-style-type: none"> • * – Match all rpc-names.
notification-name	<p>Name of the notification node this rule applies to. Optional.</p> <ul style="list-style-type: none"> • * – Match all notification-names.
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>
<config>
  <nacm 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>
  </edit-config>
  <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.

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
!
```

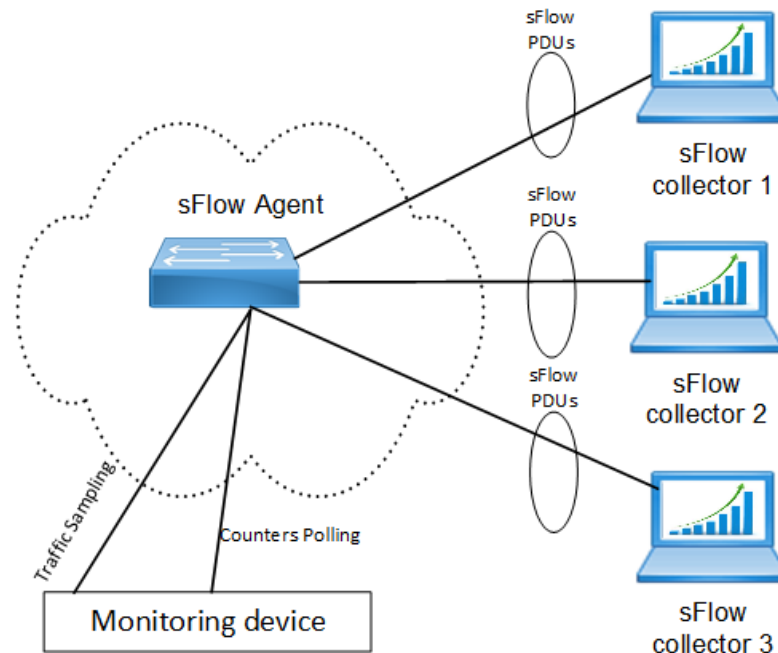
Configuration

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 8. 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-	
		Rate	Count	Interval	Count	Size (bytes)	
Ingress	Egress	Ingress	Egress	(sec)	Ingress	Egress	
xe12 256	3 16	1024	2000	3	6	10	0
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 (s) (bytes) (sec)	Collector ID	Packet-Sampling Maximum Header Rate Direction	Packet-Sampling Sampling		Packet-Sampling Interval		Counter-	
			Count		Count		Size	
			Ingress	Egress	Ingress	Egress		
xe12 120	1	1029 egress-only	1029	0	0	0	0	
xe13 140	5	1048 ingress-only	0	0	0	0	0	
xe14 128	3	1048 ingress-only	1048	0	0	20	1248	
xe15 120	4	1029 ingress-only	0	0	0	0	0	
xe16 140	2	2048 egress-only	3020	0	0	0	0	

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          Size
(s)        Rate      Direction      Count            Interval      Count      Size
(bytes)    Direction
(sec)      Ingress   Egress         Ingress         Egress
-----
-----
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   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 ID	Packet-Sampling Maximum Header	Packet-Sampling Sampling	Packet-Sampling Interval	Counter-Count	Counter-Size
(s)	Rate	Direction	Count	Interval	Count	Size
(bytes)						
(sec)		Ingress	Egress	Ingress	Egress	
		Ingress	Egress			
xe1	1,2,3,4,5	1024	1024	0	0	20
256	128	both				0
xe20	1,2,3,4,5	1024	1024	0	0	20
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 Status      Sample Rate      Counter-Polling
          ID(s)   Ingress  Egress  Ingress  Egress  Interval(sec)
-----
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 Packet-Sampling Packet-Sampling Counter-
Polling      ID      Maximum Header  Sampling      Interval      Count      Size
(s)          (bytes)      Rate           Direction      Count          Interval      Count      Size
(sec)              Ingress      Egress      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 Polling	Collector ID	Packet-Sampling		Packet-Sampling		Counter-	
		Maximum Header Rate Direction	Sampling Count	Interval		Count	Size
				Ingress	Egress		
				Ingress	Egress		
xe1 0	4,5 128	0 egress-only	1024	0	2446	10	182
xe2 0	4,5 128	0 egress-only	2048	0	752	0	0

```
#show sflow statistics
```

sFlow Port Statistics:

Interface	Collector ID(s)	Packet-Sampling		Counter-Polling Count
		Count	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 \(page 213\)](#) - 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 \(page 214\)](#) - 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 xe12
(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, 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.

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, 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, 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.
- **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.
- Supports QoS queuing buffer utilization per interface for protocol counters.

Encoding


- gNMI proto encoding is not supported for ACL and BGP xpathes.
- 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 9. 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.

System Limits and Counters Commands

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 \(page 244\)](#) 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 \(page 245\)](#) 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 \(page 247\)](#) 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 \(page 249\)](#) command output includes the **Aggregator UP-Time** field, which displays the total UP duration for the aggregated interface.

show interface brief

The [show interface \(page 249\)](#) command output includes the **UP-Time** field, which displays the time duration for which the interface has remained in the UP state.

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)
```

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.

Default

None

Command Mode

Execution mode and Privileged execution mode

Applicability

Introduced before OcNOS version 1.3.

Example

Displays all routes in the IP routing table database.

```
#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   *> 1.1.1.0/24 is directly connected, eth1, 00:13:39
C   *> 2.2.2.0/24 is directly connected, eth3, 00:13:04
B   *> 88.88.1.2/32 [12/0] via 2.2.2.2, eth3, 00:06:37
C   *> 127.0.0.0/8 is directly connected, lo, 00:22:15
B   *> 150.1.0.0/16 [120/0] is a summary, Null, 00:11:19
B   *> 150.1.1.0/24 [200/0] via 1.1.1.1, eth1, 00:11:19
C   *> 192.168.52.0/24 is directly connected, eth0, 00:22:13

Gateway of last resort is not set
```

```
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
    runts 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 2. 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

Table 2. show interface output details (continued)

Field	Description
	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., 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.
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.

Key Term or Acronym	Description
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.
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.

Layer 3 Sub-interface	256
Overview	256
Configuration	257
Validation	259
Layer 3 Sub-interface Commands	262
Implementation Examples	266
Troubleshooting	267
Glossary	268

Layer 3 Sub-interface

Overview

A single physical interface when required to handle multiple VLAN traffic, can be divided into multiple logical interfaces called sub-interfaces.

All sub-interfaces under a physical port will use their parent port for transmitting and receiving data.

Sub-interfaces can be used for various purposes, as for inter-vlan routing to happen when router has only one physical interface, two sub-interfaces each with different IP network can be created under it and data can be routed between them.

Sub-interfaces let you divide a physical interface into multiple logical interfaces that are tagged with different VLAN identifiers. Because VLANs allow you to keep traffic separate on a given physical interface, you can increase the number of interfaces available to your network without adding additional physical interfaces.

Feature Characteristics

- Each subinterface is treated as a separate Layer 3 entity with its own IP address, routing table entries, and configuration.
- Subinterfaces are associated with VLAN IDs (via IEEE 802.1Q tagging), enabling traffic separation on the same physical link.

Benefits

- Reduces the need for multiple physical interfaces.
- Enables multiple IP subnets/VLANs over a single physical link.
- Allows flexible routing between VLANs without external Layer 3 devices.

Limitations

Queuing service policy-maps are not supported on Layer 3 sub-interfaces.

Configuration

Topology

- [Figure 10](#) shows an example of sub-interface configuration. In this example, there are two routers, R1 and R2, and the eth1 interface of R1 is connected directly to eth2 of R2 using an Ethernet cable.

Figure 10. Subinterface connections

The eth1.10 subinterface is created on R1, and eth2.10 is created on R2.



Note: Layer 3 Subinterfaces can be created on physical and LAG interfaces.

Creating a Sub-interface

Create and configure a Layer 3 sub-interface on a physical interface as follows:



Note: Before configuration meet all [Layer 3 Sub-interface \(page 256\)](#).

1. Create the sub-interface.

```
#configure terminal
(config)#interface eth1
(config-if)#interface eth1.10
```

2. Configure VLAN encapsulation and assign IP address.

```
(config-if)#encapsulation dot1q 10
(config-if)#ip address 10.10.10.1/24
```

3. Commit and exit the configuration.

```
(config-if)#commit
(config-if)#exit
```

Creating a Sub-interface with Encapsulation

Sub-interfaces can be configured with encapsulation to define how VLAN tags are handled. There are two types of encapsulation supported:

- Single Encapsulation (dot1q) – Standard VLAN tagging (IEEE 802.1Q)
- Double Encapsulation (Q-in-Q) – VLAN stacking using dot1q or dot1ad (IEEE 802.1ad)

Create and configure a Layer 3 sub-interface with encapsulation on a physical interface as follows:

1. Create the sub-interface with double encapsulation as dot1q.

```
#configure terminal
(config)#interface eth1.1010
(config-if)# encapsulation dot1q 10 inner-dot1q 10
(config-if)#ip address 192.168.1.50/24
(config-if)#commit
(config-if)#exit
```

2. Create the sub-interface with double encapsulation as dot1ad.

```
#configure terminal
(config)#interface eth1.20
(config-if)# encapsulation dot1ad 20 inner-dot1q 20
(config-if)#ip address 192.168.2.50/24
(config-if)#commit
(config-if)#exit
```

**Notes:**

- Use dot1ad ethertype (0x8100 | 0x88a8 | 0x9100 | 0x9200) command to configure the service-tpid value on parent port of a sub-interface. By this the tpid used for service tag for a sub-interface may be inherited from the one applied to parent interface.
- For any dot1ad sub-interface to be functional, dot1ad ethertype should be set to desired value as 0x88a8/0x9100/0x9200. Default value is 0x8100. To verify the ethertype value for the interface use show interface <subinterface> command.

Validation

In OcNOS, sub-interfaces appear as any physical interface in the `show running-config` or the `show ip interface brief` output and can be configured as any other interface.

The following examples display subinterface information from various `show` commands.

Note: The below command output is just for reference and is not directly related to the configuration provided above.

show interface brief

```
RTR1#show interface brief

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,      PD(Min-links) - Protocol Down
Min-links
DV - DDM Violation, NA - Not Applicable
NOM - No operational members, PVID - Port Vlan-id
HD - ESI Hold Timer Down

-----
Ethernet  Type  PVID  Mode                Status Reason  Speed Port
Interface                                     Ch #
-----
ce49      ETH   --    routed              up      none   100g  --

-----
Interface      Type                Status Reason  Speed
-----
ce49.2         SUBINTERFACE        up      --     0
ce49.3         SUBINTERFACE        up      --     0
ce49.4         SUBINTERFACE        up      --     0
ce49.5         SUBINTERFACE        up      --     0
ce49.6         SUBINTERFACE        up      --     0
```

show ip interface brief

```
RTR1#show ip interface brief

'*' - address is assigned by dhcp client

Interface      IP-Address      Admin-Status      Link-Status
ce49           unassigned      up                 up
ce49.2         49.49.2.1       up                 up
ce49.3         49.49.3.1       up                 up
ce49.4         49.49.4.1       up                 up
ce49.5         49.49.5.1       up                 up
ce49.6         49.49.6.1       up                 up
```

show ip ospf neighbor with VRF enabled

```
RTR1#show ip ospf neighbor

Total number of full neighbors: 2
OSPF process 1 VRF(default):
Neighbor ID    Pri  State                Dead Time  Address          Interface        Instance ID
```

```

4.4.4.4      1    Full/DR      00:00:32    48.48.2.2    vlan1.2      0
4.4.4.4      1    Full/DR      00:00:38    48.48.3.2    vlan1.3      0

Total number of full neighbors: 1
OSPF process 2 VRF(CUST-2):
Neighbor ID  Pri  State      Dead Time   Address      Interface    Instance ID
11.11.2.1    1    Full/DR    00:00:39    49.49.2.2    ce49.2      0

Total number of full neighbors: 1
OSPF process 3 VRF(CUST-3):
Neighbor ID  Pri  State      Dead Time   Address      Interface    Instance ID
11.11.3.1    1    Full/Backup 00:00:33    49.49.3.2    ce49.3      0

Total number of full neighbors: 1
OSPF process 4 VRF(CUST-4):
Neighbor ID  Pri  State      Dead Time   Address      Interface    Instance ID
11.11.4.1    1    Full/Backup 00:00:31    49.49.4.2    ce49.4      0

Total number of full neighbors: 1
OSPF process 5 VRF(CUST-5):
Neighbor ID  Pri  State      Dead Time   Address      Interface    Instance ID
11.11.5.1    1    Full/Backup 00:00:39    49.49.5.2    ce49.5      0

```

show ip route with VRF enabled

```

RTR1#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"
C      1.2.200.0/24 is directly connected, xe1.200, 01:29:19
O      4.4.4.4/32 [110/11] via 48.48.3.2, vlan1.3, 00:37:17
       [110/11] via 48.48.2.2, vlan1.2
O      44.44.44.0/24 [110/2] via 48.48.3.2, vlan1.3, 00:37:17
       [110/2] via 48.48.2.2, vlan1.2
C      47.47.2.0/24 is directly connected, xe47.2, 00:34:42
C      48.48.2.0/24 is directly connected, vlan1.2, 00:41:19
C      48.48.3.0/24 is directly connected, vlan1.3, 00:41:19
C      127.0.0.0/8 is directly connected, lo, 01:30:09

Gateway of last resort is not set

RTR1#show ip route vrf all
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      1.2.200.0/24 is directly connected, xe1.200, 01:29:32
O      4.4.4.4/32 [110/11] via 48.48.3.2, vlan1.3, 00:37:30
       [110/11] via 48.48.2.2, vlan1.2
O      44.44.44.0/24 [110/2] via 48.48.3.2, vlan1.3, 00:37:30
       [110/2] via 48.48.2.2, vlan1.2
C      47.47.2.0/24 is directly connected, xe47.2, 00:34:55

```

```
C          48.48.2.0/24 is directly connected, vlan1.2, 00:41:32
C          48.48.3.0/24 is directly connected, vlan1.3, 00:41:32
C          127.0.0.0/8 is directly connected, lo, 01:30:22
IP Route Table for VRF "management"
C          127.0.0.0/8 is directly connected, lo.management, 01:30:22
C          192.168.10.0/24 is directly connected, eth0, 01:30:22
IP Route Table for VRF "CUST-1"
C          127.0.0.0/8 is directly connected, lo.CUST-1, 01:30:22
IP Route Table for VRF "CUST-2"
C          1.1.2.0/24 is directly connected, xe1.2, 01:29:35
C          1.2.101.0/24 is directly connected, xe1.101, 01:29:34
C          1.3.201.0/24 is directly connected, xe1.201, 01:29:32
O          11.11.2.0/24 [110/20] via 49.49.2.2, ce49.2, 00:51:06
O          11.12.101.0/24 [110/20] via 49.49.2.2, ce49.2, 00:51:06
O          11.13.201.0/24 [110/20] via 49.49.2.2, ce49.2, 00:51:06
C          49.49.2.0/24 is directly connected, ce49.2, 01:29:31
C          127.0.0.0/8 is directly connected, lo.CUST-2, 01:30:22
IP Route Table for VRF "CUST-3"
C          1.1.3.0/24 is directly connected, xe1.3, 01:29:35
C          1.2.102.0/24 is directly connected, xe1.102, 01:29:34
C          1.3.202.0/24 is directly connected, xe1.202, 01:29:32
O          11.11.3.0/24 [110/20] via 49.49.3.2, ce49.3, 01:12:44
O          11.12.102.0/24 [110/20] via 49.49.3.2, ce49.3, 01:12:44
O          11.13.202.0/24 [110/20] via 49.49.3.2, ce49.3, 01:12:44
C          49.49.3.0/24 is directly connected, ce49.3, 01:29:31
```

Layer 3 Sub-interface Commands

Below are the commands for Layer 3 sub-interface:

encapsulation

Use this command to configure encapsulation under a sub-interface. Using this command, a Layer 3 sub-interface can be configured as a port-vlan or stacked vlan. Before configuring the encapsulation on sub-interface, the operating state of the sub-interface is admin down. After configuring the encapsulation, the operating state of the sub-interface becomes up.

Command Syntax

```
encapsulation ((dot1q|dot1ad) VLAN_ID (inner-dot1q VLAN_ID|))  
no encapsulation
```

Parameters

dot1q

IEEE 802.1Q VLAN-tagged packets

dot1ad

IEEE 802.1ad VLAN-tagged packets

VLAN_ID

First (outer) VLAN identifier on the sub-interface. The outer VLAN ID configured on a sub-interface must be within the valid VLAN range supported by the platform (1–4094), excluding the VLAN reserved on the parent interface.

inner-dot1q VLAN_ID

Second (inner 802.1Q) VLAN identifier on the sub-interface.

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 3.0. For DC - OcNOS version 7.0.0 is the applicability.

Example

```
#configure terminal  
(config-if)#interface xe9.1  
(config-if)#encapsulation dot1q 10  
  
(config-if)#interface xe9.2  
(config-if)#encapsulation dot1q 20 inner-dot1q 20  
  
(config-if)#interface xe9.3  
(config-if)#encapsulation dot1ad 30  
  
(config-if)#interface xe9.4  
(config-if)#encapsulation dot1ad 40 inner-dot1q 40
```

interface IFNAME.SUBINTERFACE_ID

Use this command to configure a sub-interface. Sub-interfaces are supported for raw Ethernet interfaces as well as dynamic/static LAG interfaces. The maximum number of sub-interfaces that can be created on a parent port is platform-specific. For example, on a single physical port, up to 4060 sub-interfaces can be created on TD3, while up to 4030 sub-interfaces can be created on TH2.

Use `no` form of this command to unconfigure a sub-interface.

Command Syntax

```
interface IFNAME
no interface IFNAME
```

Parameters

IFNAME

Specifies the sub-interface name, represented in the format interface-name.sub-interface ID (for example, xe1.5, where xe1 is the interface name and 5 is the sub-interface ID). The valid sub-interface range is <1-8192>

Command Mode

Interface mode

Applicability

This command was introduced in OcNOS version 3.0. For DC - OcNOS version 7.0.0 is the applicability.

Example

```
(config)#int xe1.5
(config-if)#exit
(config)#no interface xe1.5

(config)#int po1
(config-if)#exit
(config)#int po1.1
(config-if)#exit
(config)#no interface po1.1

(config)#int sa1
(config-if)#exit
(config)#int sa1.1
(config-if)#exit
(config)#no interface sa1.1
```

show interface IFNAME.SUBINTERFACE_ID

Use this command to display the details of the sub-interface. This command displays the information about the operating status, hardware address, VRF binding details, and input/output counters. This command can display details of the a sub-interface for a dynamic/static LAG as well.

Command Syntax

```
show interface IFNAME
```

Parameters

IFNAME

Specifies the sub-interface name, represented in the format interface-name.sub-interface ID (for example, xe1.5, where xe1 is the interface name and 5 is the sub-interface ID). The valid sub-interface range is <1-8192>

Command Mode

Execution mode

Applicability

This command was introduced in OcNOS version 3.0. For DC - OcNOS version 7.0.0 is the applicability.

Example

```
#show int xe1.1
Interface xe1.1
  Hardware is SUBINTERFACE Current HW addr: 6cb9.c500.1647
  Physical:6cb9.c500.1647 Logical:(not set)
  Port Mode is Router
  Interface index: 20482049
  Metric 1 mtu 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: Never
  inet6 fe80::6eb9:c5ff:fe00:1647/64
  RX
    unicast packets 0 multicast packets 0 broadcast packets 0
    input packets 0 bytes 0
    jumbo packets 0
    undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
    input error 0
    input with dribble 0 input discard 0
    Rx pause 0
  TX
    unicast packets 0 multicast packets 0 broadcast packets 0
    output packets 0 bytes 0
    jumbo packets 0
    output errors 0 collision 0 deferred 0 late collision 0
    output discard 0
    Tx pause 0

#show int sa1.1
Interface sa1.1
```

```

Hardware is SUBINTERFACE Current HW addr: 6cb9.c500.1647
Physical:6cb9.c500.1647 Logical:(not set)
Port Mode is Router
Interface index: 409602049
Metric 1 mtu 1500
<UP,BROADCAST,RUNNING,MULTICAST>
VRF Binding: Not bound
Label switching is disabled
No Virtual Circuit configured
Administrative Group(s): None
DHCP client is disabled.
Last Flapped: Never
Statistics last cleared: Never
inet6 fe80::6eb9:c5ff:fe00:1647/64
RX
  unicast packets 0 multicast packets 0 broadcast packets 0
  input packets 0 bytes 0
  jumbo packets 0
  undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
  input error 0
  input with dribble 0 input discard 0
  Rx pause 0
TX
  unicast packets 0 multicast packets 0 broadcast packets 0
  output packets 0 bytes 0
  jumbo packets 0
  output errors 0 collision 0 deferred 0 late collision 0
  output discard 0
  Tx pause 0

#show int pol.3
Interface pol.3
  Hardware is SUBINTERFACE Current HW addr: 0030.abf1.0ec8
  Physical:0030.abf1.0ec8 Logical:(not set)
  Port Mode is Router
  Interface index: 204802051
  Metric 1 mtu 1500
  <UP,BROADCAST,RUNNING,MULTICAST>
  VRF Binding: Not bound
  Label switching is disabled
  No Virtual Circuit configured
  Administrative Group(s): None
  DHCP client is disabled.
  Last Flapped: Never
  Statistics last cleared: Never
  inet 23.0.0.2/24 broadcast 23.0.0.255
  inet6 fe80::230:abff:fef1:ec8/64
RX
  unicast packets 0 multicast packets 0 broadcast packets 0
  input packets 141805 bytes 9643544
  jumbo packets 0
  undersize 0 oversize 0 CRC 0 fragments 0 jabbers 0
  input error 0
  input with dribble 0 input discard 0
  Rx pause 0
TX
  unicast packets 0 multicast packets 0 broadcast packets 0
  output packets 0 bytes 0
  jumbo packets 0
  output errors 0 collision 0 deferred 0 late collision 0
  output discard 0
  Tx pause 0

```

Implementation Examples

Here is an example scenario and a solution for implementing L3 sub-interface.

Scenario 1: A router has only one physical interface but needs to route traffic between two different IP networks/VLANs.

Use Case 1: Create two sub-interfaces under the same physical port, each assigned to a different VLAN and IP subnet.

Scenario 2: A service provider needs to support up to 2,000 unique customers or services on a single high-speed physical link.

Use Case 2: Divide a physical interface (e.g., eth1) into multiple logical sub-interfaces (e.g., eth1.1 through eth1.2000).

Scenario 3: A provider needs to transport multiple customer VLANs over a single service provider VLAN.

Use Case 3: Configure a sub-interface with Double Encapsulation (dot1q or dot1ad). Example Command: `encapsulation dot1q 10 inner-dot1q 10`.

Troubleshooting

1. Sub-interface is in an 'Admin Down' or 'Protocol Down' state.

- Check if encapsulation is configured. Before encapsulation is applied, the sub-interface operating state remains admin down.
- Ensure the `encapsulation dot1q` or `encapsulation dot1ad` command is present in the `show running-config`.
- Confirm the physical or LAG parent interface is operationally up; sub-interfaces depend on the parent port for all data transmission.
- For `dot1ad` sub-interfaces, verify the `dot1ad ethertype` (e.g., `0x88a8`, `0x9100`) is correctly set on the parent interface.

2. Traffic is not passing through a configured Sub-interface.

- Check the interface status and assigned IP using `show ip interface brief`.
- Ensure the sub-interface IP subnet appears as a "connected" route in the routing table via `show ip route`.
- Confirm the `VLAN_ID` (outer) and `inner-dot1q VLAN_ID` (for double-tagged) match the expected incoming traffic tags.
- If routing protocols are used, verify neighbors are reaching the Full state on the sub-interface using `show ip ospf neighbor`.

Glossary

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

Key Terms/Acronym	Description
AC	Attachment Circuit: A physical or logical interface connecting customer-facing services to a Provider Edge (PE) router.
CE	Customer Edge: A customer-owned device connected to a provider's PE router via an Attachment Circuit.
FRR	Fast Reroute: An MPLS Traffic Engineering technique that provides sub-50 ms protection by creating pre-sigaled detour/bypass LSPs used when a link or node fails.
IGP	Interior Gateway Protocol: A routing protocol used within a single autonomous system. Examples in this document include IS-IS and OSPF.
IS-IS	Intermediate System to Intermediate System: An Interior Gateway Protocol (IGP) that floods link state information throughout a network of routers. Each IS-IS router independently builds a database of the network's topology, aggregating the flooded network information. A Routing Information Base (RIB) is calculated from the database by constructing a shortest path tree (SPT).
LSP	Label Switched Path: A sequence of routers that cooperatively perform Multi-Protocol Label Switching (MPLS) operations for a packet stream. An LSP is a unidirectional, point-to-point, half-duplex connection carrying information downstream from the ingress (first) router to the egress (last) router. The ingress and egress routers cannot be the same device.
PSN	The MPLS/IP core network that interconnects PE routers and transports encapsulated services.
VC	Virtual Circuit: A logical path used to transport service traffic between two PE routers over the PSN. Supports redundancy with primary/secondary configurations.