



Interface and Hardware Component Configuration Guide for Cisco 8000 Series Routers, IOS XR Release 24.1.x, 24.2.x, 24.3.x

First Published: 2024-03-14

Last Modified: 2024-09-04

Americas Headquarters

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

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

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

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

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

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

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

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

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

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

© 2024 Cisco Systems, Inc. All rights reserved.



CONTENTS

PREFACE

Preface xvii

Changes to This Document xvii

Communications, Services, and Additional Information xvii

CHAPTER 1

New and Changed Feature Information 1

Interface and Hardware Component Features Added or Modified in IOS XR Release 24.x.x 1

CHAPTER 2

YANG Data Models for Interfaces and Hardware Component Features 3

Using YANG Data Models 3

CHAPTER 3

Preconfiguring Physical Interfaces 5

Prerequisites for Preconfiguring Physical Interfaces 5

Information About Preconfiguring Physical Interfaces 6

Physical Interface Preconfiguration Overview 6

Benefits of Interface Preconfiguration 6

Use of the Interface Preconfigure Command 6

Active and Standby RPs and Virtual Interface Configuration 7

How to Preconfigure Physical Interfaces 7

CHAPTER 4

Advanced Configuration and Modification of the Management Ethernet Interface 9

Prerequisites for Configuring Management Ethernet Interfaces 9

Information About Configuring Management Ethernet Interfaces 10

Default Interface Settings 10

How to Perform Advanced Management Ethernet Interface Configuration 10

Configure a Management Ethernet Interface 10

Verify Management Ethernet Interface Configuration 13

Configuration Examples for Management Ethernet Interfaces	13
Configuring a Management Ethernet Interface: Example	13

CHAPTER 5**Configuring Ethernet Interfaces 15**

Prerequisites for Configuring Ethernet Interfaces	16
Information About Configuring Ethernet	16
Cisco 8000 Modular Line Cards	16
Default Configuration Values for 100-Gigabit Ethernet	16
Layer 2 VPN on Ethernet Interfaces	17
Gigabit Ethernet Protocol Standards Overview	17
IEEE 802.3 Physical Ethernet Infrastructure	18
IEEE 802.3ae 10-Gbps Ethernet	18
IEEE 802.3ba 100 Gbps Ethernet	18
MAC Address	18
Ethernet MTU	18
IP MTU	19
IP MTU Checks	20
IP MTU Configuration Guidelines	21
IP MTU Limitations and Feature Support	23
IP MTU Scale	24
Configure IP MTU	24
Flow Control on Ethernet Interfaces	26
802.1Q VLAN	27
Subinterfaces on the Router	27
Layer 2, Layer 3, and EFPs	30
Untagged L2 Subinterface	32
Enhanced Performance Monitoring for Layer 2 Subinterfaces (EFPs)	35
Other Performance Management Enhancements	35
Frequency Synchronization and SyncE	36
LLDP	36
LLDP Frame Format	37
LLDP TLV Format	38
Specifying User-Defined LLDP TLV Values	38
LLDP Operation	39

Supported LLDP Functions	40
Unsupported LLDP Functions	41
Setting the Carrier Delay on Physical Interfaces	41
Guidelines and Restrictions for Setting the Carrier Delay on Physical Interfaces	42
Configure the Carrier-delay Timer	43
How to Configure Ethernet	43
Configuring LLDP	44
LLDP Default Configuration	44
Enabling LLDP Per Interface	44
Enabling LLDP Globally	45
Configuring Global LLDP Operational Characteristics	46
Disabling Transmission of Optional LLDP TLVs	48
Disabling LLDP Receive and Transmit Operation for an Interface	49
Verifying the LLDP Configuration	50
Verifying the LLDP Global Configuration	51
Verifying the LLDP Interface Configuration	51
Configuring LLDP Snoop	52
Configuration Examples for Ethernet	57
Configuring an Ethernet Interface: Example	57
Configuring LLDP: Examples	58
Configuring a Layer 2 VPN AC: Example	58
Configuring Physical Ethernet Interfaces	58
Viewing Interface Counters Report	62
Instant Display of Traffic Rates for all the Physical Interfaces	63
How to Configure Interfaces in Breakout Mode	64
Information About Breakout	64
Configure Breakout in a Port	64
Remove the Breakout Configuration	64
Verify a Breakout Configuration	65
Ethernet Interface Route Statistics	65
CHAPTER 6	Configuring Ethernet OAM
	69
Information About Configuring Ethernet OAM	69
Ethernet Link OAM	70

- Neighbor Discovery 70
- EFD 71
- MIB Retrieval 71
- Miswiring Detection (Cisco-Proprietary) 72
- SNMP Traps 72
- Configuration Examples for Ethernet OAM 72
 - Configuring Ethernet OAM Features on an Individual Interface: Example 72
 - Configuring an Ethernet OAM Profile Globally: Example 72
 - Configuring Ethernet OAM Features to Override the Profile on an Individual Interface: Example 73
 - Clearing Ethernet OAM Statistics on an Interface: Example 73
 - Enabling SNMP Server Traps on a Router: Example 74
- Ethernet CFM 74
 - Maintenance Domains 76
 - Services 78
 - Maintenance Points 78
 - MEP and CFM Processing Overview 79
 - CFM Protocol Messages 80
 - Continuity Check (IEEE 802.1ag and ITU-T Y.1731) 80
 - Loopback (IEEE 802.1ag and ITU-T Y.1731) 83
 - Linktrace (IEEE 802.1ag and ITU-T Y.1731) 84
 - Configurable Logging 86
- How to Configure Ethernet OAM 86
 - Configuring Ethernet OAM 86
 - Configuring an Ethernet OAM Profile 86
 - Attaching an Ethernet OAM Profile to an Interface 91
 - Configuring Ethernet OAM at an Interface and Overriding the Profile Configuration 92
 - Verifying the Ethernet OAM Configuration 93
 - Configuring Ethernet CFM 94
 - Configuring a CFM Maintenance Domain 94
 - Configuring Services for a CFM Maintenance Domain 96
 - Enabling and Configuring Continuity Check for a CFM Service 97
 - Configuring Cross-Check on a MEP for a CFM Service 99
 - Configuring Other Options for a CFM Service 101
 - Configuring CFM MEPs 103

Configuring Y.1731 AIS	105
Verifying the CFM Configuration	108
CFM Over Bundles	108
Ethernet Frame Delay Measurement for L2VPN Services	109
Link Loss Forwarding	113
Link State Monitor and Propagation by CFM	114
Restrictions for Link Loss Forwarding for CFM	114
Configure Link Loss Forwarding for CFM	115
Remote Link State Propagation	118
Configure Link Loss Forwarding for Layer 2 Transport	119

CHAPTER 7**IP Event Dampening 121**

IP Event Dampening Overview	121
Interface State Change Events	122
Suppress Threshold	122
Half-Life Period	122
Reuse Threshold	122
Maximum Suppress Time	123
Affected Components	123
Route Types	123
Supported Protocols	123
How to Configure IP Event Dampening	124
Enabling IP Event Dampening	124
Verifying IP Event Dampening	124

CHAPTER 8**Configuring Link Bundling 125**

Limitations and Compatible Characteristics of Ethernet Link Bundles	126
Prerequisites for Configuring Link Bundling on a Router	127
Information About Configuring Link Bundling	127
Link Bundling Overview	128
Link Aggregation Through LACP	128
IEEE 802.3ad Standard	129
Configuring LACP Fallback	129
LACP Short Period Time Intervals	130

Load Balancing	131
Layer 3 Egress Load Balancing on Link Bundles	131
Configuring the Default LACP Short Period Time Interval	132
Configuring Custom LACP Short Period Time Intervals	133
QoS and Link Bundling	135
Link Bundle Configuration Overview	135
Nonstop Forwarding During Card Failover	135
Link Failover	136
Link Switchover	136
LACP Fallback	136
Designate a Member Link as Unviable	137
Guidelines and Restrictions for Designating Member Links as Unviable	138
How to Configure Link Bundling	138
Configuring Ethernet Link Bundles	138
Configuring VLAN Bundles	142
143	
VLANs on an Ethernet Link Bundle	146
Configuring a Member Link as Unviable	146
Configuration Examples for Link Bundling	147
Example: Configuring an Ethernet Link Bundle	147
Example: Configuring a VLAN Link Bundle	149

CHAPTER 9**Configuring Traffic Mirroring 151**

Introduction to Traffic Mirroring	152
Implementing Traffic Mirroring on the Cisco 8000 Series Routers	153
ERSPAN	153
Traffic Mirroring Terminology	156
Characteristics of the Source Port	157
Characteristics of the Monitor Session	157
Supported Traffic Mirroring Types	158
ACL-Based Traffic Mirroring	158
ERSPAN over GRE IPv6	158
Configuring Partial Packet Capture Ability for ERSPAN (RX)	160
ERSPAN Traffic to a Destination in a Non-Default VRF	161

Restrictions for Traffic Mirroring	161
Configuring Traffic Mirroring	163
Configuring ACLs for Traffic Mirroring	163
Troubleshooting ACL-Based Traffic Mirroring	165
Flexible CLI for ERSPAN	165
Attaching the Configurable Source Interface	166
Introduction to ERSPAN Rate Limit	168
Topology	168
Configure ERSPAN Rate Limit	169
Introduction to Local SPAN	170
Local SPAN Overview	170
Local SPAN Supported Capabilities	170
Local SPAN Restrictions	171
Configuring Local SPAN	171
Local SPAN with ACL	172
Configuring Local SPAN with ACL	172
Local SPAN Rate Limit	173
Traffic Mirroring with DSCP	174
DSCP Marking on Egress GRE Tunnel in ERSPAN	175
Configure DSCP Marking on Egress GRE Tunnel in ERSPAN	175
DSCP Bitmask to Filter Ingress ERSPAN Traffic	177
Configure DSCP Bitmask to Filter Ingress ERSPAN Traffic	177
Monitor Multiple ERSPAN Sessions with SPAN and Security ACL	179
Configure Multiple Monitor ERSPAN Sessions with SPAN and Security ACL	179
SPAN to File	180
Action Commands for SPAN to File	182
Configuring SPAN to File	183
Configuring SPAN to File for Truncation and Direction	185
Mirroring Forward-Drop Packets	186
Configuring Forward-Drop	187
Mirroring Buffer Drop Packets	189
Benefits of Mirroring Buffer Drop Packets	189
Guidelines and Restrictions for Mirroring Buffer Drop Packets	190
Configure Mirroring Buffer Drop Packets	190

Introduction to File Mirroring	193
Limitations	193
Configure File Mirroring	194
Traffic Mirroring Configuration Examples	194
Viewing Monitor Session Status: Example	194
Monitor Session Statistics: Example	195
Layer 3 ACL-Based Traffic Mirroring: Example	196
Troubleshooting Traffic Mirroring	197

CHAPTER 10

Configuring Virtual Loopback and Null Interfaces	203
Prerequisites for Configuring Virtual Interfaces	203
Information About Configuring Virtual Interfaces	203
Virtual Loopback Interface Overview	204
Null Interface Overview	204
Virtual Management Interface Overview	204
Active and Standby RPs and Virtual Interface Configuration	205
How to Configure Virtual Interfaces	205
Configuring Virtual Loopback Interfaces	205
Configuring Null Interfaces	206
Configuring Virtual IPv4 Interfaces	206
Configuration Examples for Virtual Interfaces	207
Configuring a Loopback Interface: Example	207
Configuring a Null Interface: Example	208
Configuring a Virtual IPv4 Interface: Example	208

CHAPTER 11

Configure GRE Tunnels	209
GRE Tunnels	209
Supported Features on a GRE Tunnel	211
Limitations for Configuring GRE Tunnels	212
Configure GRE Tunnels	213
Unidirectional GRE Encapsulation (GREv4)	214
Unidirectional GRE Decapsulation (GREv4)	214
ECMP and LAG Hashing for NVGRE Flows	216

CHAPTER 12	Configuring 802.1Q VLAN Interfaces	219
	Prerequisites for Configuring 802.1Q VLAN Interfaces	219
	Information About Configuring 802.1Q VLAN Interfaces	220
	802.1Q VLAN Overview	220
	Subinterfaces	220
	Subinterface MTU	220
	Native VLAN	221
	Layer 2 VPN on VLANs	221
	How to Configure 802.1Q VLAN Interfaces	221
	Configuring 802.1Q VLAN Subinterfaces	221
	Configuring an Attachment Circuit on a VLAN	223
	Removing an 802.1Q VLAN Subinterface	225
	Configuration Examples for VLAN Interfaces	226
	VLAN Subinterfaces: Example	226

CHAPTER 13	Configure IP-in-IP Tunnels	229
	Controlling the TTL Value of Inner Payload Header	232
	IP-in-IP Decapsulation	232
	Decapsulation Using Tunnel Source Direct	236
	Guidelines and Limitations	236
	Configure Decapsulation Using Tunnel Source Direct	237
	Configure Tunnel Destination with an Object Group	238
	ECMP Hashing Support for Load Balancing	241
	User-Defined Fields for ECMP Hashing	243
	Configure User-Defined Fields for ECMP Hashing	244

CHAPTER 14	Configuring Generic UDP Encapsulation	247
	Understand Generic UDP Encapsulation	248
	Restrictions	250
	Configure GUE	251
	Outer IP Header-Driven Hash Computation for Incoming GUE Packets	256
	Configure Outer IP Header-Driven Hash Computation for Incoming GUE Packets	257
	Flexible Assignment of UDP Port Numbers for Decapsulation	258

Guidelines for Setting up Decapsulation Using Flexible Port Numbers 258
 Restrictions 259
 Configuring Port Numbers for Decapsulation 259
 Verification 265

CHAPTER 15

Configuring 400G Digital Coherent Optics 267

Configuring Frequency 274
 Configuring Chromatic Dispersion 276
 Configuring Optical Transmit Power 278
 Configuring Muxponder Mode 280
 Configuring Modulation 282
 Configuring DAC Rate 284
 Configuring FEC 286
 Configuring Loopback 287
 Disable Auto-Squelching 288
 Configuring Performance Monitoring 289
 Configuring PM Parameters 291
 Configuring Alarms Threshold 294
 Configuring FEC Alarm Threshold 297
 Guidelines and Restrictions for Setting the FEC Alarm Thresholds 299
 Configuration Examples to Set FEC Alarm Threshold 300
 Configuring FDD Alarm Thresholds 300
 Configuring FED Alarm Thresholds 301
 Media Link-down PreFEC Degrade Enablement 302
 Configure Media Link-down PreFEC Degrade 303
 Alarms Troubleshooting 305
 CD Alarm 305
 Clear the CD Alarm 305
 DGD Alarm 306
 Clear the DGD Alarm 306
 FLEXO_LOF 306
 Clear the FLEXO_LOF Alarm 306
 FLEXO_LOM 306
 Clear the FLEXO_LOM Alarm 307

HI-LASERBIAS Alarm	307
Clear the HI-LASERBIAS Alarm	307
HI-RXPOWER Alarm	307
Clear the HI-RXPOWER Alarm	308
HI-RXPOWER Warn	308
Clear the HI-RXPOWER Warn Alarm	308
HI-TEMP Alarm	308
Clear the HI-TEMP Alarm	308
HI-TEMP Warn	309
Clear the HI-TEMP Warn Alarm	309
HI-TXPOWER Alarm	309
Clear the HI-TXPOWER Alarm	309
HI-TXPOWER Warn	310
Clear the HI-TXPOWER Warn Alarm	310
IMPROPER-REM	310
Clear the IMPROPER-REM Alarm	310
LOF	310
Clear the LOF Alarm	311
LOL	311
Clear the LOL Alarm	311
LOM	311
Clear the LOM Alarm	311
LO-RXPOWER Alarm	312
Clear the LO-RXPOWER Alarm	312
LO-RXPOWER Warn	312
Clear the LO-RXPOWER Warn Alarm	312
LOS	313
Clear the LOS Alarm	313
LOS-P	313
Clear the LOS-P Alarm	313
LO-TXPOWER Alarm	313
Clear the LO-TXPOWER Alarm	314
LO-TXPOWER Warn	314
Clear the LO-TXPOWER Warn Alarm	314

OOR_CD	314
Clear the OOR_CD Alarm	314
OSNR Alarm	315
Clear the OSNR Alarm	315
UNC-WORD Alarm	315
Clear the UNC-WORD Alarm	315
WVL-OOL	316
Clear the WVL-OOL Alarm	316

CHAPTER 16**Configuring Controllers 317**

How to Configure Controllers	318
Configuring Optics Controller	318
Disabling Optical Modules	318
Diagnostic Parameters for Optical Transceivers	321
View Coherent Optical Transceiver Module Parameters	326
Display of Alarms for Coherent Optical Transceivers	329
Restrictions	330
Verification	330
Loopback on Optical Transceivers	332
Media Side Input Loopback Configuration	334
Media Side Output Loopback	335
Host Side Input Loopback Configuration	336
Host Side Output Loopback Configuration	338

CHAPTER 17**Configuring QDD Optical Line System 341**

Overview of QDD Optical Line System	342
Supported Wavelength or Frequency Configuration	344
Functional Description of QDD OLS	345
QDD OLS Configurations	346
Configuring the Operational Mode, Amplifier Gain, and Amplifier Output Power	346
Configuring the Low-Threshold Power	349
Configuring the Optical Safety Remote Interlock (OSRI)	351
Configuring Safety Control Mode	352
Configuring Force APR	354

Use Case for QDD OLS pluggable	356
8-Channel Optical Line System	356
OLS Alarms Troubleshooting	357



Preface

This guide describes the interface and hardware component configuration details for Cisco IOS XR software. This chapter contains details on the changes made to this document.

- [Changes to This Document, on page xvii](#)
- [Communications, Services, and Additional Information, on page xvii](#)

Changes to This Document

This table lists the technical changes made to this document since it was first released.

Table 1: Changes to This Document

Date	Summary
September 2024	Republished for Release 24.3.1
June 2024	Republished for Release 24.2.11
March 2024	Initial release of this document

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at [Cisco Profile Manager](#).
- To get the business impact you're looking for with the technologies that matter, visit [Cisco Services](#).
- To submit a service request, visit [Cisco Support](#).
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit [Cisco DevNet](#).
- To obtain general networking, training, and certification titles, visit [Cisco Press](#).
- To find warranty information for a specific product or product family, access [Cisco Warranty Finder](#).

Cisco Bug Search Tool

[Cisco Bug Search Tool](#) (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.



CHAPTER 1

New and Changed Feature Information

This table summarizes the new and changed feature information for the *Interfaces Configuration Guide for Cisco 8000 Series Routers* for Cisco 8000 Series Routers, and tells you where they are documented.

- [Interface and Hardware Component Features Added or Modified in IOS XR Release 24.x.x](#), on page 1

Interface and Hardware Component Features Added or Modified in IOS XR Release 24.x.x

This table summarizes the new and changed feature information for the *Interfaces Configuration Guide for Cisco 8000 Series Routers* for Cisco 8000 Series Routers, and tells you where they are documented.

Table 2: New and Changed Features

Feature	Description	Introduced in Release	Where Documented
Remote Link State Propagation	This feature was introduced.	Release 24.3.1	Remote Link State Propagation , on page 118
Media Link-down Prefec Degrade Enablement	This feature was introduced.	Release 24.3.1	Media Link-down PreFEC Degrade Enablement
Untagged L2 Subinterface	This feature was introduced.	Release 24.2.11	Untagged L2 Subinterface , on page 32
Default Carrier Delay on Physical Interfaces	This feature was introduced.	Release 24.2.11	Setting the Carrier Delay on Physical Interfaces
Mirroring Buffer Drop Packets	This feature was introduced.	Release 24.2.11	Mirroring Buffer Drop Packets
Monitoring Layer 3 Connectivity Using Down MEP on L3 Interfaces	This feature was introduced.	Release 24.2.11	Ethernet CFM
User-Defined Fields for ECMP Hashing	This feature was introduced.	Release 24.2.11	User-Defined Fields for ECMP Hashing

Feature	Description	Introduced in Release	Where Documented
QDD Optical Line System	This feature was introduced.	Release 24.1.1	Configuring QDD Optical Line System, on page 341
Configurable FDD and FED Alarm Threshold Values	This feature was introduced.	Release 24.1.1	Configuring FEC Alarm Threshold, on page 297
View Additional Alarms for Coherent Optical Transceiver	This feature was introduced.	Release 24.1.1	Display of Alarms for Coherent Optical Transceivers, on page 329
View Coherent Optical Transceiver Module Parameters	This feature was introduced.	Release 24.1.1	View Coherent Optical Transceiver Module Parameters , on page 326
Generic UDP Decapsulation for IPv6 Traffic	This feature was introduced.	Release 24.1.1	Configure GUE
Extended Support for DP04QSDD-HE0 Optical Module	This feature was introduced.	Release 24.1.1	Configuring 400G Digital Coherent Optics
Enhanced Alarm Prioritization, Monitoring, and Management	This feature was introduced.	Release 24.1.1	Alarms Troubleshooting



CHAPTER 2

YANG Data Models for Interfaces and Hardware Component Features

This chapter provides information about the YANG data models for Interface and Hardware Component features.

- [Using YANG Data Models, on page 3](#)

Using YANG Data Models

Cisco IOS XR supports a programmatic way of configuring and collecting operational data of a network device using YANG data models. Although configurations using CLIs are easier and human-readable, automating the configuration using model-driven programmability results in scalability.

The data models are available in the release image, and are also published in the [Github](#) repository. Navigate to the release folder of interest to view the list of supported data models and their definitions. Each data model defines a complete and cohesive model, or augments an existing data model with additional XPath. To view a comprehensive list of the data models supported in a release, navigate to the **Available-Content.md** file in the repository.

You can also view the data model definitions using the [YANG Data Models Navigator](#) tool. This GUI-based and easy-to-use tool helps you explore the nuances of the data model and view the dependencies between various containers in the model. You can view the list of models supported across Cisco IOS XR releases and platforms, locate a specific model, view the containers and their respective lists, leaves, and leaf lists presented visually in a tree structure. This visual tree form helps you get insights into nodes that can help you automate your network.

To get started with using the data models, see the *Programmability Configuration Guide*.



CHAPTER 3

Preconfiguring Physical Interfaces

This module describes the preconfiguration of physical interfaces.

The system supports preconfiguration for the following interfaces:

- 10-Gigabit Ethernet
- 40-Gigabit Ethernet
- 100-Gigabit Ethernet
- 400-Gigabit Ethernet
- Management Ethernet

Preconfiguration allows you to configure line cards before you insert them into the router. When you insert the cards, they are instantly configured. The system creates the preconfiguration information in a different system database tree, rather than with the regularly configured interfaces. That database tree is known as the *preconfiguration directory* on the Route Processor.

There might be some preconfiguration data that you cannot verify unless the line card is present. This is because the verifiers themselves run only on the line card. You can verify such preconfiguration data when you insert the line card and initiate the verifiers. The system rejects a configuration if errors are found when you copy the configuration from the preconfiguration area to the active area.



Note Gigabit Ethernet interface is not supported. You can only preconfigure physical interfaces.

- [Prerequisites for Preconfiguring Physical Interfaces, on page 5](#)
- [Information About Preconfiguring Physical Interfaces, on page 6](#)
- [How to Preconfigure Physical Interfaces, on page 7](#)

Prerequisites for Preconfiguring Physical Interfaces

Before preconfiguring physical interfaces, ensure that you meet the following condition(s):

- Preconfiguration drivers and files are installed. Although it might be possible to preconfigure physical interfaces without a preconfiguration driver installed. The preconfiguration files are required to set the interface definition file on the router that supplies the strings for valid interface names.

Information About Preconfiguring Physical Interfaces

To preconfigure interfaces, you must understand the following concepts:

Physical Interface Preconfiguration Overview

Preconfiguration is the process of configuring interfaces before they are present in the system. You cannot verify or apply preconfigured interfaces until you insert the actual interface into the router with the matching location. The location can be the rack, slot, or module. When you insert the anticipated line card and create the interface, the system verifies the precreated configuration information. If the verification is successful, the system immediately applies the running configuration of the router.



Note When you plug the anticipated line card in, ensure that you verify any preconfiguration by using the appropriate **show** commands.

Use the **show run** command to see the interfaces that are in the preconfigured state.



Note We recommend filling out preconfiguration information in your site planning guide. This allows you to compare the anticipated configuration with the actual preconfigured interfaces when you install the card and the interfaces are up.



Tip Use the **commit best-effort** command to save the preconfiguration to the running configuration file. The **commit best-effort** command merges the target configuration with the running configuration and commits only the valid configuration (best effort). Some configuration might fail due to semantic errors, but the valid configuration still comes up.

Benefits of Interface Preconfiguration

Preconfigurations reduce downtime when you add new cards to the system. With preconfiguration, you can instantly configure the new modular services card that actively runs during the line card bootup.

Another advantage of performing a preconfiguration is that during a card replacement, when you remove the line card, you can still see the previous configuration and make modifications.

Use of the Interface Preconfigure Command

To preconfigure the interfaces that are not yet present in the system, use the **interface preconfigure** command in global configuration mode.

The **interface preconfigure** command places the router in interface configuration mode. You must be able to add any possible interface commands. The verifiers registered for the preconfigured interfaces verify the

configuration. The preconfiguration is complete when you enter the **end** command, or any matching exit or global configuration mode command.



Note It is possible that you are not able to verify some configurations until you insert the line card is inserted. Do not enter the **no shutdown** command for new preconfigured interfaces, because the no form of this command removes the existing configuration, and there is no existing configuration.

You must provide names during preconfiguration that matches with the name of the interface that is created. If the interface names do not match, the system does not apply preconfiguration when the interface is created. The interface names must begin with the interface type that is supported by the router and for which drivers have been installed. However, the slot, port, subinterface number, and channel interface number information cannot be validated.



Note Specifying an interface name that already exists and is configured (or an abbreviated name like Hu0/3/0/0) is not permitted.

Active and Standby RPs and Virtual Interface Configuration

The standby RP is available and is in a state in which it can take the load from the an active RP, if required. Following are the conditions when a standby RP becomes an active RP:

- Failure detection by a watchdog.
- Standby RP is administratively commanded to take over.
- Removal of the active RP from the chassis.

If a second RP is not present in the chassis while the first is in operation, the system may insert a second RP. The second RP then automatically becomes the standby RP. The standby RP may also be removed from the chassis with no effect on the system other than loss of RP redundancy.

After failover, the virtual interfaces become available on the standby (now active) RP. Their state and configuration is unchanged, and there is no loss of forwarding (in the case of tunnels) over the interfaces during the failover. The routers use nonstop forwarding (NSF) over tunnels through the failover of the host RP.



Note You do not need to configure anything to guarantee that the standby interface configurations are maintained.

How to Preconfigure Physical Interfaces

This task describes only the most basic preconfiguration of an interface.

```
/* Enter global configuration mode. */  
RP/0/RP0/CPU0:router:router:hostname# configure
```

```

/* Enters interface preconfiguration mode for an interface, where type specifies
the supported interface type that you want to configure and interface-path-id specifies
the location where the interface will be located in rack/slot/module/port notation. */

RP/0/RP0/CPU0:router:router(config)# interface preconfigure HundredGigE 0/3/0/2

/* Assign an IP address and mask to the interface. Use one of the following commands:
- ipv4 address ip-address subnet-mask
- ipv4 address ip-address/prefix */
RP/0/RP0/CPU0:router(config-if-pre)# ipv4 address 192.168.1.2/31
RP/0/RP0/CPU0:router(config-if-pre)# end
or
RP/0/RP0/CPU0:router(config-if-pre)# commit
RP/0/RP0/CPU0:router# show running-config

```

- When you issue the **end** command, the system prompts you to commit changes: `Uncommitted changes found, commit them before exiting (yes/no/cancel)?`
- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit best-effort** command to save the configuration changes to the running configuration file and remain within the configuration session. The **commit best-effort** command merges the target configuration with the running configuration and commits only valid changes (best effort). Some configuration changes might fail due to semantic errors.



CHAPTER 4

Advanced Configuration and Modification of the Management Ethernet Interface

This module describes the configuration of Management Ethernet interfaces.

Before you use Telnet to access the router through the LAN IP address, you must set up a Management Ethernet interface and enable the Telnet servers.



Note By default, the Management Ethernet interfaces are present on the system. However, you must configure these interfaces to:

- Access the router.
- Use protocols and applications, such as Simple Network Management Protocol (SNMP), HTTP, eXtensible Markup Language (XML), TFTP, Telnet, and Command-Line Interface (CLI.)

-
- [Prerequisites for Configuring Management Ethernet Interfaces, on page 9](#)
 - [Information About Configuring Management Ethernet Interfaces, on page 10](#)
 - [How to Perform Advanced Management Ethernet Interface Configuration, on page 10](#)
 - [Configuration Examples for Management Ethernet Interfaces, on page 13](#)

Prerequisites for Configuring Management Ethernet Interfaces

Before you perform the Management Ethernet interface configuration procedures that are described in this chapter, ensure that you meet the following tasks and conditions:

- You have performed the initial configuration of the Management Ethernet interface.
- You know how to apply the generalized interface name specification *rack/slot/module/port*.



Note For transparent switchover, ensure that both the active and standby Management Ethernet interfaces are physically connected to the same LAN or switch.

Information About Configuring Management Ethernet Interfaces

To configure Management Ethernet interfaces, you must understand the following concept(s):

Default Interface Settings

This table describes the default Management Ethernet interface settings that you can change with manual configuration. The system does not display the default settings in the **show running-config** command output.

Table 3: Management Ethernet Interface Default Settings

Parameter	Default Value	Configuration File Entry
Speed in Mbps	Default speed is 1G with autonegotiated.	Speed is non-configurable.
Duplex mode	Default duplex mode is full-duplex with autonegotiated.	Duplex mode is non-configurable.
MAC address	MAC address is read from the hardware burned-in address (BIA).	MAC address is non-configurable.

How to Perform Advanced Management Ethernet Interface Configuration

This section contains the following procedures:

Configure a Management Ethernet Interface

Perform this task to configure a Management Ethernet interface. This procedure provides the minimal configuration that is required for the Management Ethernet interface.



Note The maximum MTU value for the management interface MgmtEth0/RP0/CPU0/0 is 9678 bytes.

```
RP/0/RP0/CPU0:router # configure

/* Enter interface configuration mode and specify the Ethernet interface name and notation
  rack/slot/module/port. */
RP/0/RP0/CPU0:router (config) # interface MgmtEth 0/RP0/CPU0/0

RP/0/RP0/CPU0:router (config-if) # ipv4 address 1.76.18.150/16 (or)
ipv4 address 1.76.18.150 255.255.0.0
```

Assigns an IP address and subnet mask to the interface.

- Replace *ip-address* with the primary IPv4 address for the interface.

- Replace *mask* with the mask for the associated IP subnet. You can specify the network mask in either of the two ways:
- The network mask can be a four-part dotted decimal address. For example, 255.255.0.0 indicates that each bit equal to 1 means that the corresponding address bit belongs to the network address.
- The system indicates the network mask as a slash (/) and number. For example, /16 indicates that the first 16 bits of the mask are ones, and the corresponding bits of the address are the network address.

```
RP/0/RP0/CPU0:router(config-if)# mtu 1488
```



Note (Optional) The maximum transmission unit (MTU) value for the management interface is 9678 bytes.

- The default is 1514 bytes.
- The range for the Management Ethernet interface Interface **mtu** values is from 64 through 9678 bytes.

```
/* Remove the shutdown configuration, which removes the forced administrative down on the
interface, enabling it to move to an up or down state. */
RP/0/RP0/CPU0:router(config-if)# no shutdown
RP/0/RP0/CPU0:router(config-if)# end
or
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

```
RP/0/RP0/CPU0:router# show interfaces MgmtEth 0/RP0/CPU0/0
```

This example displays advanced configuration and verification of the Management Ethernet interface on the RP:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface MgmtEth 0/RP0/CPU0/0
RP/0/RP0/CPU0:router(config)# ipv4 address 1.76.18.150/16
RP/0/RP0/CPU0:router(config-if)# no shutdown
RP/0/RP0/CPU0:router(config-if)# commit
```

```
RP/0/RP0/CPU0:router:Mar 26 01:09:28.685 :ifmgr[190]:%LINK-3-UPDOWN :Interface
MgmtEth0/RP0/CPU0/0, changed state to Up
RP/0/RP0/CPU0:router(config-if)# end
```

```
RP/0/RP0/CPU0:router# show interfaces MgmtEth 0/RP0/CPU0/0
```

```
MgmtEth0/RP0/CPU0/0 is up, line protocol is up
Interface state transitions: 3
Hardware is Management Ethernet, address is 1005.cad8.4354 (bia 1005.cad8.4354)
Internet address is 1.76.18.150/16
MTU 1488 bytes, BW 1000000 Kbit (Max: 1000000 Kbit)
    reliability 255/255, txload 0/255, rxload 0/255
Encapsulation ARPA,
Full-duplex, 1000Mb/s, 1000BASE-T, link type is autonegotiation
loopback not set,
Last link flapped 00:00:59
ARP type ARPA, ARP timeout 04:00:00
Last input 00:00:00, output 00:00:02
Last clearing of "show interface" counters never
5 minute input rate 4000 bits/sec, 3 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
    21826 packets input, 4987886 bytes, 0 total input drops
    0 drops for unrecognized upper-level protocol
    Received 12450 broadcast packets, 8800 multicast packets
        0 runts, 0 giants, 0 throttles, 0 parity
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    1192 packets output, 217483 bytes, 0 total output drops
    Output 0 broadcast packets, 0 multicast packets
    0 output errors, 0 underruns, 0 applique, 0 resets
    0 output buffer failures, 0 output buffers swapped out
    3 carrier transitions
```

```
RP/0/RP0/CPU0:router# show running-config interface MgmtEth 0/RP0/CPU0/0
```

```
interface MgmtEth0/RP0/CPU0/0
mtu 1488
ipv4 address 1.76.18.150/16
ipv6 address 2002::14c:125a/64
ipv6 enable
!
```

The following example displays VRF configuration and verification of the Management Ethernet interface on the RP with the source address:

```
RP/0/RP0/CPU0:router# show run interface MgmtEth 0/RP0/CPU0/0
interface MgmtEth0/RP0/CPU0/0
vrf httpupload
ipv4 address 10.8.67.20 255.255.0.0
ipv6 address 2001:10:8:67::20/48
!
```

```
RP/0/RP0/CPU0:router# show run http
Wed Jan 30 14:58:53.458 UTC
http client vrf httpupload
http client source-interface ipv4 MgmtEth0/RP0/CPU0/0
```

```
RP/0/RP0/CPU0:router# show run vrf
Wed Jan 30 14:59:00.014 UTC
vrf httpupload
!
```

Verify Management Ethernet Interface Configuration

Perform this task to verify configuration modifications on the Management Ethernet interfaces.

```
RP/0/RP0/CPU0:router# show interfaces MgmtEth 0/RP0/CPU0/0
RP/0/RP0/CPU0:router# show running-config interface MgmtEth 0/RP0/CPU0/0
```

Configuration Examples for Management Ethernet Interfaces

This section provides the following configuration examples:

Configuring a Management Ethernet Interface: Example

This example displays advanced configuration and verification of the Management Ethernet interface on the RP:

```
RP/0//CPU0:router# configure
RP/0//CPU0:router(config)# interface MgmtEth 0/RP0/CPU0/0
RP/0//CPU0:router(config)# ipv4 address 172.29.52.70 255.255.255.0
RP/0//CPU0:router(config-if)# no shutdown
RP/0//CPU0:router(config-if)# commit
RP/0//CPU0:Mar 26 01:09:28.685 :ifmgr[190]:%LINK-3-UPDOWN :Interface MgmtEth 0/RP0/CPU0/0,
  changed state to Up
RP/0//CPU0:router(config-if)# end

RP/0//CPU0:router# show interfaces MgmtEth 0/RP0/CPU0/0

MMgmtEth0//CPU0/0 is up, line protocol is up
  Hardware is Management Ethernet, address is 0011.93ef.e8ea (bia 0011.93ef.e8ea
)
  Description: Connected to Lab LAN
  Internet address is 172.29.52.70/24
  MTU 1514 bytes, BW 100000 Kbit
    reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set,
  ARP type ARPA, ARP timeout 04:00:00
  Last clearing of "show interface" counters never
  5 minute input rate 3000 bits/sec, 7 packets/sec
  5 minute output rate 0 bits/sec, 1 packets/sec
    30445 packets input, 1839328 bytes, 64 total input drops
    0 drops for unrecognized upper-level protocol
  Received 23564 broadcast packets, 0 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  171672 packets output, 8029024 bytes, 0 total output drops
  Output 16 broadcast packets, 0 multicast packets
  0 output errors, 0 underruns, 0 applique, 0 resets
  0 output buffer failures, 0 output buffers swapped out
  1 carrier transitions

RP/0//CPU0:router# show running-config interface MgmtEth 0/

interface MgmtEth0/RP0/CPU0/0
  description Connected to Lab LAN
```

```
ipv4 address 172.29.52.70 255.255.255.0  
!
```




CHAPTER 5

Configuring Ethernet Interfaces

This module describes the configuration of Ethernet interfaces.

The distributed 10-Gigabit, 25-Gigabit Ethernet, 40-Gigabit, 100-Gigabit Ethernet, architecture and features deliver network scalability and performance, while enabling service providers to offer high-density, high-bandwidth networking solutions designed to interconnect the router with other systems in POPs, including core and edge routers, Layer 2 switches and Layer 3 switches.

Table 4: Feature History Table

Feature Name	Release Information	Feature Description
Introduction of IP MTU on Q200-based Systems	Release 7.5.2	You can configure IP MTU for IPv4 and IPv6 on a Layer 3 interface. Depending on your specific network requirements, this ability to specify IP MTU settings helps optimize router data transmission. Use the <code>show ipv4/ipv6 interfaces</code> command to view the IP MTU configurations.



Tip You can programmatically configure and manage the Ethernet interfaces using `openconfig-if-ethernet.yang` and `openconfig-interfaces.yang` OpenConfig data models. To get started with using data models, see the *Programmability Configuration Guide for Cisco 8000 Series Routers*.

- [Prerequisites for Configuring Ethernet Interfaces, on page 16](#)
- [Information About Configuring Ethernet, on page 16](#)
- [Setting the Carrier Delay on Physical Interfaces, on page 41](#)
- [How to Configure Ethernet, on page 43](#)
- [Viewing Interface Counters Report, on page 62](#)
- [How to Configure Interfaces in Breakout Mode, on page 64](#)
- [Ethernet Interface Route Statistics, on page 65](#)

Prerequisites for Configuring Ethernet Interfaces

Before configuring Ethernet interfaces, ensure that you meet the following conditions:

- Access to Cisco 8200 series routers or Cisco 8800 series routers with at least one of the supported line cards installed.
- Know the interface IP address.
- Starting with Cisco IOS XR Release 24.1.1, the `openconfig-if-ip.yang` open config model supports a new leaf, `oc-if:interfaces/oc-if:interface/oc-if:subinterfaces/oc-if:subinterface/ipv4/addresses/address[ip]/config/type`. When using the `openconfig-if-ip.yang` model to update or revise the IP address for an interface, you must specify the IP address as either primary or secondary. You can have only one IP address as primary, but you can have multiple addresses configured as secondary.
- Ensure to specify the generalized interface name with the standard notation of `rack/slot/module/port`.

Information About Configuring Ethernet

This section provides the following information:

Cisco 8000 Modular Line Cards

The current release of the Cisco 8800 Series Routers support the following line cards:

- 36-port QSFP56-DD 400 GbE Line Card - This line card provides 14.4 Tbps via 36 QSFP56-DD ports. It also supports 100G, 2x100G, and 400G modules. If 36 of 2x100G modules are used, the line card can have 72 HundredGigE interfaces.
- 48-port QSFP28 100 GbE Line Card - This line card provides 4.8 Tbps with MACsec support on all ports. It also supports QSFP+ optics for 40G compatibility.

The 8800 Series line cards utilize multiple #ChipName forwarding ASICs to achieve high performance and bandwidth with line rate forwarding.

Default Configuration Values for 100-Gigabit Ethernet

This table describes the default interface configuration parameters that are present when an interface is enabled on a 36-port Line Card or a 48-port Line Card.



Note You must use the **shutdown** command to bring an interface administratively down. The interface default is **no shutdown**. When a line card is first inserted into the router, if there is no established preconfiguration for it, the configuration manager adds a shutdown item to its configuration. This shutdown can be removed only by entering the **no shutdown** command.

Table 5: 100-Gigabit Ethernet Line Card Default Configuration Values

Parameter	Configuration File Entry	Default Value
Flow control	flow-control	egress off ingress off
MTU	mtu	<ul style="list-style-type: none"> • 1514 bytes for normal frames • 1518 bytes for 802.1Q tagged frames. • 1522 bytes for Q-in-Q frames.
MAC address	mac address	Hardware burned-in address (BIA)

Layer 2 VPN on Ethernet Interfaces

Layer 2 Virtual Private Network (L2VPN) connections emulate the behavior of a LAN across an L2 switched, IP or MPLS-enabled IP network, allowing Ethernet devices to communicate with each other as if they were connected to a common LAN segment.

The L2VPN feature enables service providers (SPs) to provide Layer 2 services to geographically disparate customer sites. Typically, an SP uses an access network to connect the customer to the core network. On the router, this access network is typically Ethernet.

Traffic from the customer travels over this link to the edge of the SP core network. The traffic then tunnels through an L2VPN over the SP core network to another edge router. The edge router sends the traffic down another attachment circuit (AC) to the customer's remote site.

On the router, an AC is an interface that is attached to an L2VPN component, such as a bridge domain.

The L2VPN feature enables users to implement different types of end-to-end services.

Switching takes place through local switching where traffic arriving on one AC is immediately sent out of another AC without passing through a pseudowire.

Keep the following in mind when configuring L2VPN on an Ethernet interface:

- L2VPN links support QoS (Quality of Service) and MTU (maximum transmission unit) configuration.
- If your network requires that packets are transported transparently, you may need to modify the packet's destination MAC (Media Access Control) address at the edge of the Service Provider (SP) network. This prevents the packet from being consumed by the devices in the SP network.

Use the **show interfaces** command to display AC information.

To attach Layer 2 service policies, such as QoS, to the Ethernet interface, refer to the appropriate Cisco IOS XR software configuration guide.

Gigabit Ethernet Protocol Standards Overview

The Gigabit Ethernet interfaces support the following protocol standards:

These standards are further described in the sections that follow.

IEEE 802.3 Physical Ethernet Infrastructure

The IEEE 802.3 protocol standards define the physical layer and MAC sublayer of the data link layer of wired Ethernet. IEEE 802.3 uses Carrier Sense Multiple Access with Collision Detection (CSMA/CD) access at various speeds over various physical media. The IEEE 802.3 standard covers 10 Mbps Ethernet. Extensions to the IEEE 802.3 standard specify implementations for 40-Gigabit Ethernet and 100-Gigabit Ethernet.

IEEE 802.3ae 10-Gbps Ethernet

Under the International Standards Organization's Open Systems Interconnection (OSI) model, Ethernet is fundamentally a Layer 2 protocol. 10-Gigabit Ethernet uses the IEEE 802.3 Ethernet MAC protocol, the IEEE 802.3 Ethernet frame format, and the minimum and maximum IEEE 802.3 frame size. 10-Gbps Ethernet conforms to the IEEE 802.3ae protocol standards.

Just as 1000BASE-X and 1000BASE-T (Gigabit Ethernet) remained true to the Ethernet model, 10-Gigabit Ethernet continues the natural evolution of Ethernet in speed and distance. Because it is a full-duplex only and fiber-only technology, it does not need the carrier-sensing multiple-access with the CSMA/CD protocol that defines slower, half-duplex Ethernet technologies. In every other respect, 10-Gigabit Ethernet remains true to the original Ethernet model.

IEEE 802.3ba 100 Gbps Ethernet

IEEE 802.3ba is supported on the Cisco 1-Port 100-Gigabit Ethernet PLIM beginning in Cisco IOS XR 7.0.11.

MAC Address

A MAC address is a unique 6-byte address that identifies the interface at Layer 2.

Ethernet MTU

The Ethernet maximum transmission unit (MTU) is the size of the largest frame, minus the 4-byte frame check sequence (FCS), that the system transmits on the Ethernet network. Every physical network along the destination of a packet can have a different MTU.

Cisco IOS XR software supports two types of frame forwarding processes:

- Fragmentation for IPv4 packets – In this process, IPv4 packets are fragmented as necessary to fit within the MTU of the next-hop physical network.



Note IPv6 does not support fragmentation.

- MTU discovery process determines largest packet size – This process is available for all IPv6 devices, and for originating IPv4 devices. In this process, the originating IP device determines the size of the largest IPv6 or IPv4 unfragmented packet that the system can send. The largest packet is equal to the smallest MTU of any network between the IP source and the IP destination devices. If a packet is larger than the smallest MTU of all the networks in its path, the system fragments that packet as necessary. This process ensures that the originating device does not send an IP packet that is too large.

The system automatically enables the jumbo frame support for frames that exceed the standard frame size. The default value is 1514 for standard frames and 1518 for 802.1Q tagged frames. These numbers exclude the 4-byte frame check sequence (FCS).

IP MTU

In IP protocol, Maximum Transmission Unit (MTU) refers to the maximum size of an IP packet that the system transmits without fragmentation over a given medium. The size of an IP packet includes IP headers but excludes headers from the data link layer, also known as the Ethernet headers. The default IP MTU on all router interfaces is 1500 bytes, when IP is enabled by using the IP address configuration commands. However, you can configure the IP MTU to different value as well.

Starting Cisco IOS XR Release 7.5.2, the system supports IP MTU (IPv4 and IPv6) on Q200-based systems on the following Cisco 8000 Series router and line card:

- 8201-32FH
- 88-LC0-36FH-M

How is Ethernet MTU different from IP MTU?

Ethernet MTU defines the maximum packet size that an interface supports, while IP MTU defines the MTU size of an IP packet.

How is IP MTU calculated?

The following scenarios provide information on how the system calculates IP MTU size, when:

- Ethernet MTU is not configured, the system sets:
 - Default value as 1514 bytes on a physical or bundle main interface
 - 1518 bytes for single tag VLAN interface
 - 1522 bytes for double tag VLAN (QinQ) interface



Note In this case, IP MTU value will be 1500 bytes for IPv4 and IPv6 MTUs.

- Ethernet MTU is configured as X bytes, the system sets:
 - X bytes on a physical or bundle main interface
 - X+4 bytes for single tag VLAN interface
 - X+8 bytes for double tag VLAN (QinQ) interface



Note In this case, IP MTU is X-14 bytes for IPv4 and IPv6 MTUs.

The following are some of the important guidelines for IP MTU size:

- When no Ethernet MTU and IP MTU is configured, the default value is 1500B.

- When Ethernet MTU and no IP MTU is configured, IP MTU value in the hardware is Ethernet MTU-14B.
- IP MTU value can't be more than Ethernet MTU value. For example, if the Ethernet MTU value is 3000 bytes and you configure IP MTU value as 5000 bytes. The system sets the IP MTU value as 2986 (3000-14)

What is Maximum Receive Unit (MRU) and how is it different from MTU?

MRU is the largest packet size that an interface can receive. This is an ingress parameter. Usually, MRU equals MTU. However, you can't configure an MRU value. The Ethernet MTU, also known as a Layer2 (L2) value that you configure on a physical interface is also applied as the MRU of that physical interface.

The following table lists Ethernet MTU, IPv4, and IPv6 MTU support across various platforms and their limitations as applicable:

Table 6: IP MTU Support Across Platforms

ASIC	Ethernet MTU Check	IPv4 and IPv6 MTU
Q100-based	Ethernet MTU check on an egress interface is not supported.	Supported and the system derives IP MTU value from Ethernet MTU.
Q200-based	Ethernet MTU check on an egress interface is not supported.	Supported Note IPv4 and IPv6 can have their own separate MTU values.

IP MTU Checks

Cisco IOS XR supports MRU checks, ethernet MTU checks, and IP MTU checks. These checks decide if an IP payload packets needs fragmentation or not. Ethernet MTU and IP MTU check is applied on all egress traffic for each packet that flows in the system. In the forwarding plane, also known as dataplane, the IP packet length is compared against IP MTU value applied on the L3 forwarding interface. Full packet length is compared against ethernet MTU value applied on the physical port.

The following table describes IP MTU check to each forwarding flow on Q200-based and Q100-based systems.

Table 7: IP MTU Check

ASIC	IP MTU Check
Q100-based	These interfaces implement IP MTU checks and provide fragmentation for all IP payload packets. A single value for both IPv4 and IPv6 is supported.
Q200-based	These interfaces implement IP MTU checks and provide fragmentation for all IP payload packets.
Mixed System with Q100 and Q200-based line cards	Interfaces with Q100-based system provides its behavior and interfaces with Q200-based systems provides its behavior as described in this table.

IP MTU Configuration Guidelines

An Ethernet interface has a default of 1514 bytes. IP MTU is always derived from the default Ethernet MTU. If Ethernet MTU is configured, IP MTU is derived from the previously configured Ethernet MTU.

The following are the configuration guidelines for IP MTU across various platforms:

Guidelines for Q200-based Systems

- If you don't define any Ethernet MTU configuration, the system assigns the default value of Ethernet packet size of 1514 bytes.
- For IPv4 and IPv6 MTUs, 1500 bytes is used, which is derived by subtracting 14 bytes of Ethernet header size from its default value.
- On any interface, if an Ethernet MTU is configured, the new configuration takes higher precedence than the default interface configuration. However, in such a scenario, the new configuration doesn't get applied on the subinterfaces. The rest of the interfaces continue to work with the default Ethernet MTU configurations.
- For the double tag (QinQ) packets, MRU and Ethernet MTU value are derived by adding 8 bytes to the configured value.
- IPv4 and IPv6 MTU values are derived by subtracting 14 bytes of the Ethernet header size out of the configured value.
- If an IPv4 or IPv6 MTU is configured, it is applied to IP MTU value. MRU or MTU values remain unaffected by this configuration.
- Configuration restrictions apply to validate that the IP MTU value is smaller than Ethernet MTU value that is configured on an interface.

The following sample table explains how the system calculates the Ethernet MTU, IPv4, and IPv6 MTU values when configured on various interfaces that are Q200-based systems.

Table 8: Interfaces and Configurations

Interface Type	Default Configurations <i>when, no Ethernet or IP MTU configuration is applied</i>	Ethernet MTU Configurations <i>for example, MTU = 1614</i>	IPv4/IPv6 Configurations <i>for example, IPv4 IPv6 MTU = 1550, Ethernet MTU = 1614</i>
Any Main Interface	NA	Configuration that is applied on any main interface and its subinterfaces is effective. Configuration that is applied on any subinterface is not effective.	NA
Physical Port	MRU = 1514 + 8	MRU = 1614 + 8	MRU = 1614 + 8

Interface Type	Default Configurations <i>when, no Ethernet or IP MTU configuration is applied</i>	Ethernet MTU Configurations <i>for example, MTU = 1614</i>	IPv4/IPv6 Configurations <i>for example, IPv4 IPv6 MTU = 1550, Ethernet MTU = 1614</i>
L3 Main Interface (Physical or bundle)	Ethernet MTU = 1514 + 8 IPv4/IPv6 MTU = 1500	Ethernet MTU = 1614 + 8 IPv4/IPv6 MTU = 1614 - 14	Ethernet MTU = 1614 + 8 IPv4/IPv6 MTU = 1550
L3 Sub-interface (Physical or bundle)	Ethernet MTU = 1514 + 8 IPv4/IPv6 MTU = 1500	Note Configuration not applicable on subinterfaces. Takes main interface configuration. Ethernet MTU = Not applicable IPv4/IPv6 MTU = 1614 - 14	Ethernet MTU = 1614 + 8 IPv4/IPv6 MTU = 1550
SVI/BVI	IPv4/IPv6 MTU = 1500	Ethernet MTU = Not applicable IPv4/IPv6 MTU = 1600	IPv4/IPv6 MTU = 1550
Ethernet Main Interface	Ethernet MTU = 1514 + 8	Ethernet MTU = 1614 + 8 IPv4/IPv6 MTU = Not applicable	Not applicable Ethernet MTU = 1614 + 8
Ethernet Sub-interface	Ethernet MTU = 1514 + 8	Ethernet MTU = 1614 + 8 IPv4/IPv6 MTU = Not applicable	Not applicable Ethernet MTU = 1614 + 8

Guidelines for Q100-based Systems

- If you don't define any MTU configuration in the system, the system assigns default value of 1514 bytes to the Ethernet packets.
- An MRU of 1522 bytes is set to allow for any double tag packets on an interface.
- Ethernet MTU check on egress interface is not supported.
- For IPv4 and IPv6 MTUs, 1500 bytes is used, which is derived by subtracting 14 bytes of Ethernet header size from its default value.
- If you configure an IP MTU (IPv4/IPv6 MTU) on any interface, the configuration does not take effect.

The following table lists the MTU configurations and MTU calculations on various interfaces for Q100-based systems:

Table 9: Interfaces and Configurations

Interface Type	Default Configurations	Interface MTU Configurations <i>where, MTU = 1614</i>
Any Main Interface	NA	Configuration applied on any main interface and its subinterfaces is effective. Configuration applied on any sub-interface is not effective.
Physical Port	MRU = 1514 + 8	MRU = 1614 + 8
L3 Main Interface (Physical or bundle)	Ethernet MTU check on egress interface is not supported. IPv4/IPv6 MTU = 1522 applied on full packet size	Ethernet MTU check on egress interface is not supported. IPv4/IPv6 MTU = 1622 applied on full packet size
L3 Subinterface (Physical or bundle)	Ethernet MTU check on egress interface is not supported. IPv4/IPv6 MTU = 1522 applied on full packet size	Note Configuration not applicable on sub-interfaces. Takes main interface configuration. Ethernet MTU = Not applicable IPv4/IPv6 MTU = 1622 applied on full packet size
SVI and BVI	Ethernet MTU is not applicable IPv4/IPv6 MTU = 1522 applied on full packet size	
Ethernet Main Interface	Ethernet MTU check on egress interface is not supported. IPv4/IPv6 MTU is not applicable	
Ethernet Subinterface	Ethernet MTU check on egress interface is not supported. IPv4/IPv6 MTU is not applicable	
GRE tunnel	Ethernet MTU check on egress interface is not supported. IPv4/IPv6 MTU = 1522 applied on full packet size	

IP MTU Limitations and Feature Support

The following are the limitations and feature support for IP MTU:

- MPLS MTU is not supported.
- GRE/IPinIP MTU is not supported.



Note GRE interface uses the configured L3 MTU. If an egress packet is above the configured L3 MTU value, the packet is discarded.

- Multicast MDT MTU is not supported.
- Ethernet MTU configuration on sub-interfaces is not supported.
- Ethernet MTU configuration on BVI interfaces is not supported.



Important Although CLIs for the unsupported features might be available on your router, they are not functional.

IP MTU Scale

The following are the scale support and limitations for IP MTU configurations:

- Ethernet MTU configuration is allowed on every physical or bundle interface without any scale limits.
- IPv4/IPv6 MTU configuration is allowed on every L3 forwarding interface. However, the system applies it as an IP MTU profile with a combination of <IPv4 MTU and IPv6 MTU> set. Each profile takes a unique set of these two MTUs, and supports eight such unique MTUs. The IP MTU profiles are stored identically across all NPUs of a given line card.
- When system runs out of resources (OOR), the system generates syslog messages while the continuing to use the previously configured IP MTU values.

Configure IP MTU

To configure IPv4, or IPv6 MTU, you must be in the configuration mode on L3 interface.

To configure IP MTU, use the following configuration steps:

```
Router#config t
Router#interface HundredGigE0/0/0/33
Router:abc(config)#
Router:abc(config-if)#ipv4 mtu 2000
Router:abc(config-subif)#commit
Router:abc(config-if)#end
```

Running Configuration

```
Router:abc#
Router:abc#sh ipv4 int HundredGigE0/0/0/33
Thu Apr 28 16:54:10.820 UTC
HundredGigE0/0/0/33 is Shutdown, ipv4 protocol is Down
  Vrf is default (vrfid 0x60000000)
  Internet address is 10.10.10.91/29
  MTU is 9642 (2000 is available to IP)
```

Verification

```

Router:abc#sh int HundredGigE0/0/0/33
Thu Apr 28 16:57:23.390 UTC
HundredGigE0/0/0/33 is administratively down, line protocol is administratively down
  Interface state transitions: 0
  Hardware is HundredGigE, address is e41f.7bde.123c (bia e41f.7bde.123c)
  Internet address is 194.19.242.94/29
  MTU 9642 bytes, BW 100000000 Kbit (Max: 100000000 Kbit)
    reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation ARPA,
  Full-duplex, 100000Mb/s, 100GBASE-AOC, link type is force-up
  output flow control is off, input flow control is off
  Carrier delay (up) is 10 msec
  loopback not set,
  ARP type ARPA, ARP timeout 04:00:00
  Last input never, output never
  Last clearing of "show interface" counters never
  30 second input rate 0 bits/sec, 0 packets/sec
  30 second output rate 0 bits/sec, 0 packets/sec
    0 packets input, 0 bytes, 0 total input drops
    0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 0 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity
    0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
    0 packets output, 0 bytes, 0 total output drops
  Output 0 broadcast packets, 0 multicast packets
    0 output errors, 0 underruns, 0 applique, 0 resets
    0 output buffer failures, 0 output buffers swapped out
    0 carrier transitions

```

```
Router:abc#
```

Configure IPv4 MTU

To configure IPv4 MTU, use the following configuration steps:

```

Router#config t
Router#interface HundredGigE0/0/0/33
Router:abc(config)#
Router:abc(config-if)#ipv4 mtu 2000
Router:abc(config-subif)#commit
Router:abc(config-if)#end

```

Running Configuration

```

Router:abc#sh run int HundredGigE0/0/0/33
Thu Apr 28 16:22:06.796 UTC
interface HundredGigE0/0/0/33
mtu 9642
ipv4 mtu 2000
ipv4 address 10.10.10.1 255.255.255.248
ipv6 address 10::1:10:9/112
load-interval 30
!

```

Verification

```
Router:abc#sh ipv4 int HundredGigE0/0/0/33
Thu Apr 28 16:54:10.820 UTC
HundredGigE0/0/0/33 is Shutdown, ipv4 protocol is Down
  Vrf is default (vrfid 0x60000000)
  Internet address is 10.10.10.1/29
  MTU is 9642 (2000 is available to IP)
```

Configure IPv6 MTU

To configure IPv6 MTU, use the following configuration steps:

```
Router#config t
Router#interface HundredGigE0/0/0/33
Router:abc(config)#
Router:abc(config-if)#ipv6 mtu 3000
Router:abc(config-subif)#commit
Router:abc(config-if)#end
```

Running Configuration

```
Router:abc#sh run int HundredGigE0/0/0/33
Thu Apr 28 16:23:09.141 UTC
interface HundredGigE0/0/0/33
mtu 9642
ipv4 mtu 2000
ipv4 address 10.10.10.1 255.255.255.248
ipv6 mtu 3000
ipv6 address 10::10:10:9/112
load-interval 30
!

Router:abc#
```

Verification

```
Router:abc#sh ipv6 int HundredGigE0/0/0/33
Thu Apr 28 16:54:41.222 UTC
HundredGigE0/0/0/33 is Shutdown, ipv6 protocol is Down, Vrfid is default (0x60000000)
  IPv6 is enabled, link-local address is fe80::e61f:7bff:fede:123c [TENTATIVE]
  Global unicast address(es):
    194::19:242:94, subnet is 10::10:10:0/112 [TENTATIVE]
  Joined group address(es): ff02::2 ff02::1
  MTU is 9642 (3000 is available to IPv6)
```

Flow Control on Ethernet Interfaces

The flow control that the system uses on 10-Gigabit Ethernet interfaces consists of periodically sending flow control pause frames. It is fundamentally different from the usual full and half-duplex flow control that is used on standard management interfaces. You can activate or deactivate flow control for ingress traffic only. The system automatically implements flow control for egress traffic.

802.1Q VLAN

A VLAN is a group of devices on one or more LANs that the system configures so that they can communicate as if they are attached to the same wire, when in fact they are located on several different LAN segments. Because VLANs are based on logical instead of physical connections, it is flexible for user and host management, bandwidth allocation, and resource optimization.

The IEEE's 802.1Q protocol standard addresses the problem of breaking large networks into smaller parts so broadcast and multicast traffic does not consume more bandwidth than necessary. The standard also helps to provide a higher level of security between segments of internal networks.

The 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames.

Subinterfaces on the Router

In Cisco IOS XR, interfaces are, by default, main interfaces. A main interface is also known as a trunk interface, which you must not confuse with the word trunk in the context of VLAN trunking.

There are two types of trunk interfaces:

- Physical
- Bundle

On the router, the system automatically creates the physical interfaces when the router recognizes a card and its physical interfaces. However, the system does not automatically create bundle interfaces but you must create them at the time of configuration.

The following configuration samples are examples of the trunk interfaces that you can create:

- interface HundredGigE 0/5/0/0
- interface bundle-ether 1

A subinterface is a logical interface that the system create under a trunk interface.

To create a subinterface, you must first identify a trunk interface under which to place it. In case of bundle interfaces, if a trunk interface does not exist, you must create a bundle interface before creating any subinterfaces under it.

You can then assign a subinterface number to the subinterface that you want to create. The subinterface number must be a positive integer from zero to some high value. For a given trunk interface, each subinterface under it must have a unique value.

Subinterface numbers do not need to be contiguous or in numeric order. For example, the following subinterfaces numbers are valid under one trunk interface:

1001, 0, 97, 96, 100000

Subinterfaces can never have the same subinterface number under one trunk.

In the following example, the card in slot 5 has trunk interface, HundredGigE 0/5/0/0. A subinterface, HundredGigE 0/5/0/0.0, is created under it.

```
RP/0/RSP0/CPU0:router# conf
Mon Sep 21 11:12:11.722 EDT
RP/0/RSP0/CPU0:router(config)# interface HundredGigE0/5/0/0.0
```

```

RP/0/RSP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RSP0/CPU0:router(config-subif)# commit

RP/0/RSP0/CPU0:Sep 21 11:12:34.819 : config[65794]: %MGBL-CONFIG-6-DB_COMMIT : Configuration
  committed by user 'root'. Use 'show configuration commit changes 1000000152' to view the
  changes.

RP/0/RSP0/CPU0:router(config-subif)# end

RP/0/RSP0/CPU0:Sep 21 11:12:35.633 : config[65794]: %MGBL-SYS-5-CONFIG_I : Configured from
  console by root
RP/0/RSP0/CPU0:router#

```

The **show run** command displays the trunk interface first, then the subinterfaces in ascending numerical order.

```

RP/0/RSP0/CPU0:router# show run | begin HundredGigE 0/5/0/0
Mon Sep 21 11:15:42.654 EDT
Building configuration...
interface HundredGigE 0/5/0/0
  shutdown
!
interface HundredGigE 0/5/0/0.0
  encapsulation dot1q 100
!
interface HundredGigE 0/5/0/1
  shutdown
!

```

When a subinterface is first created, the router recognizes it as an interface that, with few exceptions, is interchangeable with a trunk interface. After the new subinterface is configured further, the **show interface** command can display it along with its unique counters:

The following example shows the display output for the trunk interface, HundredGigE 0/5/0/0, followed by the display output for the subinterface HundredGigE 0/5/0/0.0.

```

RP/0/RSP0/CPU0:router# show interface HundredGigE 0/5/0/0
Mon Sep 21 11:12:51.068 EDT
HundredGigE0/5/0/0 is administratively down, line protocol is administratively down.
  Interface state transitions: 0
  Hardware is HundredGigE, address is 0024.f71b.0ca8 (bia 0024.f71b.0ca8)
  Internet address is Unknown
  MTU 1514 bytes, BW 1000000 Kbit
    reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation 802.1Q Virtual LAN,
  Full-duplex, 1000Mb/s, SXFD, link type is force-up
  output flow control is off, input flow control is off
  loopback not set,
  ARP type ARPA, ARP timeout 04:00:00
  Last input never, output never
  Last clearing of "show interface" counters never
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    0 packets input, 0 bytes, 0 total input drops
    0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 0 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity
  0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
  0 packets output, 0 bytes, 0 total output drops
  Output 0 broadcast packets, 0 multicast packets
  0 output errors, 0 underruns, 0 applique, 0 resets
  0 output buffer failures, 0 output buffers swapped out
  0 carrier transitions

```

```
RP/0/RSP0/CPU0:router# show interface HundredGigE0/5/0/0.0
Mon Sep 21 11:12:55.657 EDT
HundredGigE0/5/0/0.0 is administratively down, line protocol is administratively down.
  Interface state transitions: 0
  Hardware is VLAN sub-interface(s), address is 0024.f71b.0ca8
  Internet address is Unknown
  MTU 1518 bytes, BW 1000000 Kbit
    reliability 255/255, txload 0/255, rxload 0/255
  Encapsulation 802.1Q Virtual LAN, VLAN Id 100, loopback not set,
  ARP type ARPA, ARP timeout 04:00:00
  Last input never, output never
  Last clearing of "show interface" counters never
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    0 packets input, 0 bytes, 0 total input drops
    0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 0 multicast packets
  0 packets output, 0 bytes, 0 total output drops
  Output 0 broadcast packets, 0 multicast packets
```

This example shows two interfaces being created at the same time: first, the bundle trunk interface, then a subinterface attached to the trunk:

```
RP/0/RSP0/CPU0:router# conf
Mon Sep 21 10:57:31.736 EDT
RP/0/RSP0/CPU0:router(config)# interface Bundle-Ether1
RP/0/RSP0/CPU0:router(config-if)# no shut
RP/0/RSP0/CPU0:router(config-if)# interface bundle-Ether1.0
RP/0/RSP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RSP0/CPU0:router(config-subif)# commit
RP/0/RSP0/CPU0:Sep 21 10:58:15.305 : config[65794]: %MGBL-CONFIG-6-DB_COMMIT : C
onfiguration committed by user 'root'. Use 'show configuration commit changes 10
00000149' to view the changes.
RP/0/RSP0/CPU0:router# show run | begin Bundle-Ether1
Mon Sep 21 10:59:31.317 EDT
Building configuration..
interface Bundle-Ether1
!
interface Bundle-Ether1.0
  encapsulation dot1q 100
!
```

You delete a subinterface using the **no interface** command.

```
RP/0/RSP0/CPU0:router#
RP/0/RSP0/CPU0:router# show run | begin HundredGigE 0/5/0/0
Mon Sep 21 11:42:27.100 EDT
Building configuration...
interface HundredGigE 0/5/0/0
  negotiation auto
!
interface HundredGigE 0/5/0/0.0
  encapsulation dot1q 100
!
interface HundredGigE 0/5/0/1
  shutdown
!
RP/0/RSP0/CPU0:router# conf
Mon Sep 21 11:42:32.374 EDT
RP/0/RSP0/CPU0:router(config)# no interface HundredGigE 0/5/0/0.0
RP/0/RSP0/CPU0:router(config)# commit
RP/0/RSP0/CPU0:Sep 21 11:42:47.237 : config[65794]: %MGBL-CONFIG-6-DB_COMMIT : Configuration
```

```

    committed by user 'root'. Use 'show configuration commit changes 1000000159' to view the
    changes.
RP/0/RSP0/CPU0:router(config)# end
RP/0/RSP0/CPU0:Sep 21 11:42:50.278 : config[65794]: %MGBL-SYS-5-CONFIG_I : Configured from
    console by root
RP/0/RSP0/CPU0:router# show run | begin HundredGigE 0/5/0/0
Mon Sep 21 11:42:57.262 EDT
Building configuration...
interface HundredGigE 0/5/0/0
    negotiation auto
!
interface HundredGigE 0/5/0/1
    shutdown
!

```

Layer 2, Layer 3, and EFPs

On the router, a trunk interface can be either a Layer 2 or Layer 3 interface. A Layer 2 interface is configured using the **interface** command with the **l2transport** keyword. When the **l2transport** keyword is not used, the interface is a Layer 3 interface. Subinterfaces are configured as Layer 2 or Layer 3 subinterface in the same way.

A Layer 3 trunk interface or subinterface is a routed interface and can be assigned an IP address. Traffic sent on that interface is routed.

A Layer 2 trunk interface or subinterface is a switched interface and cannot be assigned an IP address. A Layer 2 interface must be connected to an L2VPN component. Once it is connected, it is called an access connection.

Subinterfaces can only be created under a Layer 3 trunk interface. Subinterfaces cannot be created under a Layer 2 trunk interface.

A Layer 3 trunk interface can have any combination of Layer 2 and Layer 3 interfaces.

The following example shows an attempt to configure a subinterface under an Layer 2 trunk and the commit errors that occur. It also shows an attempt to change the Layer 2 trunk interface to an Layer 3 interface and the errors that occur because the interface already had an IP address assigned to it.

```

RP/0/RP0/CPU0:router# config
Mon Sep 21 12:05:33.142 EDT
RP/0/RP0/CPU0:router(config)# interface HundredGigE0/5/0/0
RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.0.0.1/24
RP/0/RP0/CPU0:router(config-if)# commit
RP/0/RP0/CPU0:Sep 21 12:05:57.824 : config[65794]: %MGBL-CONFIG-6-DB_COMMIT : Configuration
    committed by user 'root'. Use 'show configuration commit changes 1000000160' to view the
    changes.
RP/0/RP0/CPU0:router(config-if)# end
RP/0/RP0/CPU0:Sep 21 12:06:01.890 : config[65794]: %MGBL-SYS-5-CONFIG_I : Configured from
    console by root
RP/0/RP0/CPU0:router# show run | begin HundredGigE0/5/0/0
Mon Sep 21 12:06:19.535 EDT
Building configuration...
interface HundredGigE0/5/0/0
    ipv4 address 10.0.0.1 255.255.255.0
    negotiation auto
!
interface HundredGigE0/5/0/1
    shutdown
!
RP/0/RP0/CPU0:router#
RP/0/RP0/CPU0:router#

```



```

RP/0/RP0/CPU0:router# conf
Mon Sep 21 12:08:07.426 EDT
RP/0/RP0/CPU0:router(config)# interface HundredGigE0/5/0/0 l2transport
RP/0/RP0/CPU0:router(config-if-l2)# commit

% Failed to commit one or more configuration items during a pseudo-atomic operation. All
changes made have been reverted. Please issue 'show configuration failed' from this session
to view the errors
RP/0/RP0/CPU0:router(config-if-l2)# no ipv4 address
RP/0/RP0/CPU0:router(config-if)# commit
RP/0/RP0/CPU0:Sep 21 12:08:33.686 : config[65794]: %MGBL-CONFIG-6-DB_COMMIT : Configuration
committed by user 'root'. Use 'show configuration commit changes 1000000161' to view the
changes.
RP/0/RP0/CPU0:router(config-if)# end
RP/0/RP0/CPU0:Sep 21 12:08:38.726 : config[65794]: %MGBL-SYS-5-CONFIG_I : Configured from
console by root
RP/0/RP0/CPU0:router#
RP/0/RP0/CPU0:router# show run interface HundredGigE0/5/0/0
Mon Sep 21 12:09:02.471 EDT
interface HundredGigE0/5/0/0
  negotiation auto
  l2transport
  !
!
RP/0/RP0/CPU0:router#
RP/0/RP0/CPU0:router# conf
Mon Sep 21 12:09:08.658 EDT
RP/0/RP0/CPU0:router(config)# interface HundredGigE0/5/0/0.0
RP/0/RP0/CPU0:router(config)# interface HundredGigE0/5/0/0.0
RP/0/RP0/CPU0:router(config-subif)# commit

% Failed to commit one or more configuration items during a pseudo-atomic operation. All
changes made have been reverted. Please issue 'show configuration failed' from this session
to view the errors
RP/0/RP0/CPU0:router(config-subif)#
RP/0/RP0/CPU0:router(config-subif)# interface HundredGigE0/5/0/0
RP/0/RP0/CPU0:router(config-if)# no l2transport
RP/0/RP0/CPU0:router(config-if)# interface HundredGigE0/5/0/0.0
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 99
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 11.0.0.1/24
RP/0/RP0/CPU0:router(config-subif)# interface HundredGigE0/5/0/0.1 l2transport
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 700
RP/0/RP0/CPU0:router(config-subif)# commit
RP/0/RP0/CPU0:Sep 21 12:11:45.896 : config[65794]: %MGBL-CONFIG-6-DB_COMMIT : Configuration
committed by user 'root'. Use 'show configuration commit changes 1000000162' to view the
changes.
RP/0/RP0/CPU0:router(config-subif)# end
RP/0/RP0/CPU0:Sep 21 12:11:50.133 : config[65794]: %MGBL-SYS-5-CONFIG_I : Configured from
console by root
RP/0/RP0/CPU0:router#
RP/0/RP0/CPU0:router# show run | b HundredGigE0/5/0/0
Mon Sep 21 12:12:00.248 EDT
Building configuration...
interface HundredGigE0/5/0/0
  negotiation auto
  !
interface HundredGigE0/5/0/0.0
  ipv4 address 11.0.0.1 255.255.255.0
  encapsulation dot1q 99
  !
interface HundredGigE0/5/0/0.1 l2transport
  encapsulation dot1q 700
  !

```

```
interface HundredGigE0/5/0/1
  shutdown
!
```

All subinterfaces must have unique encapsulation statements, so that the router can send incoming packets and frames to the correct subinterface. If a subinterface does not have an encapsulation statement, the router will not send any traffic to it.

In Cisco IOS XR, an Ethernet Flow Point (EFP) is implemented as a Layer 2 subinterface, and consequently, a Layer 2 subinterface is often called an EFP.

A Layer 2 trunk interface can be used as an access connection. However, a Layer 2 trunk interface is not an EFP because an EFP, by definition, is a substream of an overall stream of traffic.

Cisco IOS XR also has other restrictions on what can be configured as a Layer 2 or Layer 3 interface. Certain configuration blocks only accept Layer 3 and not Layer 2. For example, OSPF only accepts Layer 3 trunks and subinterface. Refer to the appropriate Cisco IOS XR configuration guide for other restrictions.

Untagged L2 Subinterface

Table 10: Feature History Table

Feature Name	Release Information	Feature Description
Untagged L2 Subinterface	Release 24.2.11	<p>You can now use untagged L2 subinterfaces to effectively manage and process traffic from customer edge (CE) devices that do not employ VLAN tagging. This capability allows you to apply services to untagged packets, which would not have been possible if the packets were to be logically received on the main interface. As a result, you can now push a dot1q or other supported Layer 2 encapsulation on the received frame.</p> <p>This feature introduces the encapsulation untagged command.</p>

Untagged L2 subinterface

An untagged L2 subinterface is a subinterface on a network device that isn't linked to a VLAN tag. The VLAN tags divide network traffic into smaller, distinct networks, which helps improve security and manageability in a bigger network setup.

Challenges in Processing Traffic without Untagged L2 Subinterfaces

You cannot perform the following actions on an untagged frame received on the main interface:

- Apply commands available only to l2transport subinterfaces, such as popping or pushing a dot1q or dot1ad header.
- Attach a service policy because you cannot simultaneously attach a service policy to the main interface and subinterfaces of that interface.

Untagged L2 Subinterface Offerings

The untagged subinterface takes a higher priority than the main interface on service mapping. When both interfaces are created, the untagged traffic is mapped to the untagged subinterface.

Untagged Traffic Mapped to an Untagged Subinterface

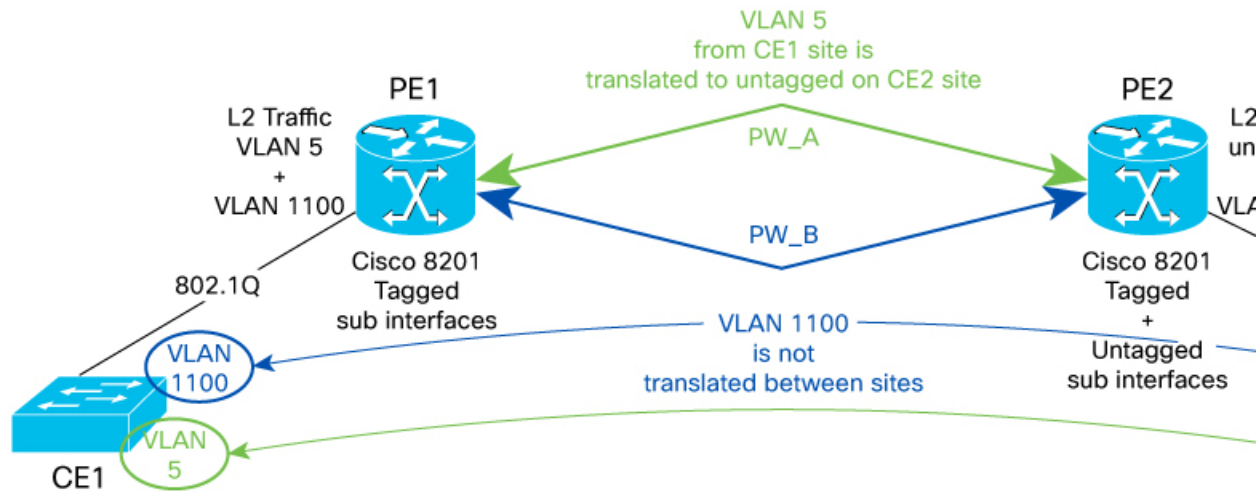
Traffic using an unsupported VLAN tag format is considered untagged traffic and mapped to an untagged subinterface. The following VLAN tag formats are supported:

- Single tag format - dot1q
- Double tag format - outer dot1ad paired with inner dot1q tag, or outer dot1q tag paired with inner dot1q tag
- Three or more tags - outermost dot1ad tag followed by dot1q tag, or outermost dot1q tag followed by dot1q tag

Benefits of Untagged L2 Subinterfaces

- Manages untagged and tagged traffic on the same port to different services.
- Allows rewriting of different VLANs on untagged subinterfaces for both tagged and untagged traffic.

How does the Untagged L2 subinterface Work?



Let's understand the feature using this sample topology, in which the P2P L2VPN service is deployed between the two CE devices.

In this topology,

- VLAN 1100 is a common VLAN that is shared between CE1 and CE2.
PE nodes do not require VLAN translation. However, if a VLAN ID is not common on CE1 and CE2, PE must perform VLAN translation.
- For instance, VLAN 5 is known only to CE1, and for the same service, CE1 uses VLAN 5, while CE2 uses no VLAN. In this situation, PE1 performs VLAN translation.

- When traffic is sent from CE1 to CE2, PE1 removes the VLAN 5 tag on the traffic from CE1 and then sends untagged traffic to PE2 through the pseudowire (PW). PE2 disposes the PW header and forwards untagged traffic to CE2.
- When a traffic is sent from CE2 to CE1, PE2 maps untagged traffic from CE2 to an L2 subinterface and sends the traffic to PE1 through the PW. PE1 disposes of the PW header and adds the VLAN 5 tag before switching the traffic to CE1.

Configure Untagged L2 Subinterface

To manage traffic efficiently, configure untagged L2 subinterface.

Procedure

Step 1 Create an untagged L2 subinterface with the **encapsulation untagged** keyword under an L2 transport main interface.

This feature is only applicable to L2 transport main and subinterfaces and does not support any other encapsulation types on the same L2 subinterface.

Example:

```
Router# configure
Router(config)# interface HundredGigE0/0/16.1500 l2transport
Router(config-subif)# encapsulation untagged
Router(config-subif)# commit
```

Step 2 Verify that the untagged encapsulation is configured on the L2 subinterface.

Example:

```
Router# show interfaces HundredGigE0/0/16.1500
HundredGigE0/0/16.1500 is up, line protocol is up
Interface state transitions: 1
Hardware is VLAN sub-interface(s), address is 0029.c201.3f0c
Internet address is 40.40.50.1/24
MTU 1522 bytes, BW 100000000 Kbit (Max: 100000000 Kbit)
reliability 255/255, txload 0/255, rxload 0/255
Encapsulation Untagged Virtual LAN,
Last link flapped 00:01:25
ARP type ARPA, ARP timeout 04:00:00
Last input never, output never
Last clearing of "show interface" counters never
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
0 packets input, 0 bytes, 0 total input drops
0 drops for unrecognized upper-level protocol
Received 0 broadcast packets, 0 multicast packets
0 packets output, 0 bytes, 0 total output drops
Output 0 broadcast packets, 0 multicast packets
```

Enhanced Performance Monitoring for Layer 2 Subinterfaces (EFPs)

Beginning in Cisco IOS XR Release 7.2.12, the router adds support for basic counters for performance monitoring on Layer 2 subinterfaces. This section provides a summary of the new support for Layer 2 interface counters.

The **interface basic-counters** keyword has been added to support a new entity for performance statistics collection and display on Layer 2 interfaces in the following commands:

- **performance-mgmt statistics interface basic-counters**
- **performance-mgmt threshold interface basic-counters**
- **performance-mgmt apply statistics interface basic-counters**
- **performance-mgmt apply threshold interface basic-counters**
- **performance-mgmt apply monitor interface basic-counters**
- show performance-mgmt monitor interface basic-counters
- show performance-mgmt statistics interface basic-counters

The **performance-mgmt threshold interface basic-counters** command supports the following attribute values for Layer 2 statistics, which also appear in the **show performance-mgmt statistics interface basic-counters** and **show performance-mgmt monitor interface basic-counters** command:

Attribute	Description
InOctets	Bytes received (64-bit)
InPackets	Packets received (64-bit)
InputQueueDrops	Input queue drops (64-bit)
InputTotalDrops	Inbound correct packets discarded (64-bit)
InputTotalErrors	Inbound incorrect packets discarded (64-bit)
OutOctets	Bytes sent (64-bit)
OutPackets	Packets sent (64-bit)
OutputQueueDrops	Output queue drops (64-bit)
OutputTotalDrops	Outband correct packets discarded (64-bit)
OutputTotalErrors	Outband incorrect packets discarded (64-bit)

Other Performance Management Enhancements

The following additional performance management enhancements are included in Cisco IOS XR Release 7.0.11:

- You can retain performance management history statistics across a process restart or route processor (RP) failover using the new **history-persistent** keyword option for the **performance-mgmt statistics interface** command.

- You can save performance management statistics to a local file using the **performance-mgmt resources dump local** command.
- You can filter performance management instances by defining a regular expression group (**performance-mgmt regular-expression** command), which includes multiple regular expression indices that specify strings to match. You apply a defined regular expression group to one or more statistics or threshold templates in the **performance-mgmt statistics interface** or **performance-mgmt thresholds interface** commands.

Frequency Synchronization and SyncE

Cisco IOS XR Software supports SyncE-capable Ethernet on the router. Frequency Synchronization enables you to distribute the precision clock signals around the network. The system injects a highly accurate timing signal into the router in the network. The timing signals use an external timing technology, such as Cesium atomic clocks, or GPS, and then pass the signals to the physical interfaces of the router. Peer routers can then recover this precision frequency from the line, and also transfer it around the network. This feature is traditionally applicable to SONET or SDH networks, but is now available on Ethernet for Cisco 8000 Series Router with Synchronous Ethernet capability. For more information, see *Cisco 8000 Series Router System Management Configuration Guide*.

LLDP

Table 11: Feature History Table

Feature Name	Release Information	Feature Description
LLDP Snooping	Release 7.3.3	<p>With this release, you can further leverage the Link Layer Discovery Protocol (LLDP) information for directly attached devices or equipment in an L2 (Layer 2) network via LLDP snoop. In order to utilize the LLDP snoop functionality, the neighbouring devices must exchange the LLDP packets with the L2 network.</p> <p>With the help of the LLDP snoop functionality, you can identify the cabling and modeling failures and isolate faults.</p> <p>To enable LLDP snoop, enable the LLDP command on an interface while the outgoing (TX) traffic is disabled.</p>

The Cisco Discovery Protocol (CDP) is a device discovery protocol that runs over Layer 2 (the Data Link layer) on all Cisco-manufactured devices (routers, bridges, access servers, and switches). CDP allows network management applications to automatically discover and learn about other Cisco devices connected to the network.

To support non-Cisco devices and to allow for interoperability between other devices, the router also supports the IEEE 802.1AB Link Layer Discovery Protocol (LLDP). LLDP is also a neighbor discovery protocol that is used for network devices to advertise information about themselves to other devices on the network. This protocol runs over the Data Link Layer, which allows two systems running different network layer protocols to learn about each other.

LLDP supports a set of attributes that it uses to learn information about neighbor devices. These attributes have a defined format known as a Type-Length-Value (TLV). LLDP supported devices can use TLVs to receive and send information to their neighbors. Details such as configuration information, device capabilities, and device identity can be advertised using this protocol.

In addition to the mandatory TLVs (Chassis ID, Port ID, End of LLDPDU, and Time-to-Live), the router also supports the following basic management TLVs, which are optional:

- Port Description
- System Name
- System Description
- System Capabilities
- Management Address

These optional TLVs are automatically sent when LLDP is active, but you can disable them as needed using the **lldp tlv-select <Optional TLV> disable** command.



Note MACsec encrypts LLDP packets by default. You can enable exceptions in the MACsec policy using the **allow lldp-in-clear** command to retain the LLDP packets unencrypted with MACsec. For more information, see *MACsec Policy Exceptions for Link Layer Discovery Protocol Packets* section in the *Configuring MACsec* chapter of the *System Security Configuration Guide for Cisco 8000 Series Routers*.



Note For LLDP to work on any bundle member, enable LLDP on the bundle main interface either globally or on the interface itself. You can then choose to disable LLDP transmission on bundle main interface by using the `lldp transmit disable` command.

You can also control LLDP transmit or receive on each bundle member interface as desired.

LLDP Frame Format

LLDP frames use the IEEE 802.3 format, which consists of the following fields:

- Destination address (6 bytes)—Uses a multicast address of 01-80-C2-00-00-0E.
- Source address (6 bytes)—MAC address of the sending device or port.
- LLDP Ethertype (2 bytes)—Uses 88-CC.
- LLDP PDU (1500 bytes)—LLDP payload consisting of TLVs.
- FCS (4 bytes)—Cyclic Redundancy Check (CRC) for error checking.

LLDP TLV Format

LLDP TLVs carry the information about neighboring devices within the LLDP PDU using the following basic format:

- TLV Header (16 bits), which includes the following fields:
 - TLV Type (7 bits)
 - TLV Information String Length (9 bits)
- TLV Information String (0 to 511 bytes)

Specifying User-Defined LLDP TLV Values

It is possible to override the system default values for some of the mandatory LLDP Type-Length-Values (TLVs) that are advertised by routers to their directly connected neighboring devices. While advertising their identity and capabilities, routers can assign user-defined meaningful names instead of autogenerated values. Using the following CLIs you can specify these user-defined values:

- Router(config)#lldp system-name *system-name*
- Router(config)#lldp system-description *system-description*
- Router(config)#lldp chassis-id-type *chassis-type*
- Router(config)#lldp chassis-id *local-chassis-id*



Note The **chassis-id** value is configurable only when the **chassis-id-type** is set as **Local**. If there is a mismatch, you encounter a configuration failed error message.

The configured values, such as the system name, system description, chassis-id, chassis-type become part of the TLV in the LLDP packets that are sent to its neighbors. Values are transmitted only to LLDP enabled interfaces to which the router is connected.

You can assign any of the following values for the `chassis-id-type`. The chassis-id-types are objects that are part of the [management information base \(MIB\)](#). Depending on the selected chassis-id-type, values are assigned to these objects, and they are advertised by the router to its neighboring devices.

chassis-id-type	Description
chassis-component	Chassis identifier based on the value of entPhysicalAlias object that is defined in IETF RFC 2737.
interface-alias	Chassis identifier based on the value of ifAlias object as defined in IETF RFC 2863.
interface-name	Chassis identifier based on the name of the interface.
local	Chassis identifier based on a locally defined value.

chassis-id-type	Description
mac-address	Chassis identifier based on the value of a unicast source address.
network-address	Chassis identifier based on a network address that is associated with a particular chassis.
port-component	Chassis identifier based on the value of entPhysicalAlias object defined in IETF RFC 2737 for a port or backplane component.



Tip You can programmatically modify default values of LLDP TLVs by using the `openconfig-lldp` OpenConfig data model. To get started with using data models, see the *Programmability Configuration Guide for Cisco 8000 Series Routers*.

Configuration Example

This example shows the configuration for the LLDP TLVs that will be advertised by routers to their directly connected neighboring devices.

```
Router(config)#lldp system-name cisco-xr
Router(config)#lldp system-description cisco-xr-edge-device
Router(config)#lldp chassis-id-type local
Router(config)#lldp chassis-id ce-device9
```

Running Configuration

```
Router#show lldp
Tue Sep 13 16:03:44.550 +0530
Global LLDP information:
Status: ACTIVE
LLDP Chassis ID: ce-device9
LLDP Chassis ID Subtype: Locally Assigned Chassis Subtype
LLDP System Name: cisco-xr
LLDP advertisements are sent every 30 seconds
LLDP hold time advertised is 120 seconds
LLDP interface reinitialisation delay is 2 seconds
```

LLDP Operation

LLDP is a one-way protocol. The basic operation of LLDP consists of a device enabled for transmit of LLDP information sending periodic advertisements of information in LLDP frames to a receiving device.

Devices are identified using a combination of the Chassis ID and Port ID TLVs to create an MSAP (MAC Service Access Point). The receiving device saves the information about a neighbor in a remote lldp cache for a certain amount of time as specified in the TTL TLV received from the neighbor, before aging and removing the information.

LLDP supports the following additional operational characteristics:

- LLDP can operate independently in transmit or receive modes. On global lldp enablement, the default mode is to operate in both transmit and receive modes.

- LLDP operates as a slow protocol with transmission speeds not greater than one frame per five seconds.
- LLDP packets are sent when the following occurs:
 - The packet update frequency specified by the **lldp timer** command is reached. The default is 30 seconds.
 - When a change in the values of the managed objects occurs from the local system's LLDP MIB.
 - When LLDP is activated on an interface (3 frames are sent upon activation similar to CDP).
- When an LLDP frame is received, the LLDP remote services and PTOPO MIBs are updated with the information in the TLVs.
- LLDP supports the following actions on these TLV characteristics:
 - Interprets a neighbor TTL value of 0 as a request to automatically purge the information of the transmitting device. These shutdown LLDPDUs are typically sent prior to a port becoming inoperable.
 - An LLDP frame with a malformed mandatory TLV is dropped.
 - A TLV with an invalid value is ignored.
 - A copy of an unknown organizationally-specific TLV is maintained if the TTL is non-zero, for later access through network management.

Supported LLDP Functions

The router supports the following LLDP functions:

- IPv4 and IPv6 management addresses—In general, both IPv4 and IPv6 addresses will be advertised if they are available, and preference is given to the address that is configured on the transmitting interface.

If the transmitting interface does not have a configured address, then the system populates the TLV with an address from another interface. The advertised LLDP IP address is implemented according to the following priority order of IP addresses for interfaces on the router:

- Locally configured address on the transmitting interface
- MgmtEth0/RSP0RP0/CPU0/0
- MgmtEth0/RSP0RP0/CPU0/1
- MgmtEth0/RSP1RP1/CPU0/0
- MgmtEth0/RSP1RP1/CPU0/1
- Loopback interfaces

There are some differences between IPv4 and IPv6 address management in LLDP:

- For IPv4, as long as the IPv4 address is configured on an interface, it can be used as an LLDP management address.
- For IPv6, after the IPv6 address is configured on an interface, the interface status must be Up and pass the DAD (Duplicate Address Detection) process before it is can be used as an LLDP management address.
- LLDP is supported for the nearest physically attached, non-tunneled neighbors.

- LLDP is supported for Ethernet interfaces, L3 subinterfaces, bundle interfaces, and L3 bundle subinterfaces.
- LLDP snoop is supported on L2 interfaces, when the incoming (RX) traffic is enabled and outgoing (TX) traffic is disabled.

Unsupported LLDP Functions

The following LLDP functions are not supported on the router:

- LLDP-MED organizationally unique extension—However, interoperability still exists between other devices that do support this extension.
- Tunneled neighbors, or neighbors more than one hop away.
- LLDP TLVs cannot be disabled on a per-interface basis; However, certain optional TLVs can be disabled globally.
- LLDP SNMP trap `lldpRemTablesChange`.

Setting the Carrier Delay on Physical Interfaces

Table 12: Feature History Table

Feature Name	Release Information	Feature Description
Default Carrier Delay Value on Physical Interfaces	Release 24.2.11	<p>We have introduced the carrier-delay up default value to ensure enough time to establish a stable hardware link state. If you haven't configured the timer, the default carrier delay automatically delays the hardware link-up notifications by 200 ms.</p> <p>Previously, we recommended that you set the carrier delay-up timer to 10 ms.</p> <p>If you want to change the delay of the interface state change notification, you can use the carrier-delay command to set a different value.</p>

Feature Name	Release Information	Feature Description
Setting the Carrier Delay on Physical Interfaces	Release 7.5.4	<p>You can configure the Ethernet interfaces to delay the processing of hardware link-down and link-up notifications. With this functionality, the interface state remains stable for the configured delay duration, even if the hardware link state fluctuates. This prevents interface flapping and improves network reliability.</p> <p>Use the following CLI command in interface configuration mode to configure the delay time:</p> <p>carrier-delay</p>

Hardware links take time to stabilize after a state change and may experience link flaps. Link flap is a condition where a physical interface frequently fluctuates between an up and a down state.

During link flaps, the network reestablishes and updates routing paths after a disruption, which leads to resource exhaustion on routers. To overcome the problem, we recommend waiting until the link state is stable before taking action.

The carrier delay introduces a delay in processing interface link-state notifications in the router to provide enough time for the interface link to stabilize.

When there is a change in the link state, the carrier-delay timer starts. If the link state goes up, the **carrier-delay up** timer starts. Similarly, when the link state goes down, the **carrier-delay down** timer starts. During this delay period, the Ethernet interface state remains unchanged even if the link is physically restored. Setting a delay timer ensures the link state is established before the interface becomes operational again and avoids unnecessary interface state changes and associated traffic rerouting.

Guidelines and Restrictions for Setting the Carrier Delay on Physical Interfaces

The following usage guidelines and restrictions are applicable for setting the carrier delay on physical interfaces:

- You can configure carrier-delay for only link-up, only link-down, or both link-up and link-down notifications.
- If the **carrier-delay down milliseconds** command is configured on a physical link that fails and cannot be recovered, link down detection time increases, and it may take longer for the routing protocols to reroute the traffic around the failed link.
- Loss of Signal (LOS) is not supported on carrier delay.
- From Release 24.2.11, the default value of carrier-delay up parameter is changed to 200 ms. To restore your original configuration, you need to configure the parameter explicitly using **carrier-delay** command.
- If not configured, the carrier-delay up parameter defaults to 200 ms and the carrier-delay down parameter to 0 ms. When carrier-delay down is not configured, the higher-layer protocols are notified immediately when a physical link state changes.

- The **carrier-delay** command overwrites the previous configuration every time you execute the command. For example, if you already have the **carrier-delay up** configured and later configure the **carrier-delay down**, the carrier-delay down overwrites the previously configured carrier-delay up. However, you can configure the carrier delay up and down concurrently using the **carrier-delay up (milliseconds) down (milliseconds)** command.

Configure the Carrier-delay Timer

Default Configuration Example

In this example, one interface is brought up to check the default value of link-up notification delay.

```
Router#configure
Router(config)#interface HundredGigE 0/0/0/0
Router(config-if)#no shutdown
Router(config-if)#commit
```

Run the **show interfaces** command to check if the carrier-delay configuration for the interface defaults to 200 ms.

```
Router#show interfaces HundredGigE 0/0/0/0 | include Carrier
Fri Mar 31 07:25:05.273 UTC
Carrier delay (up) is 200 msec
```

Configuration Example

In this example, link-up and link-down notifications are configured to be delayed by 1000 ms and 150 ms using **carrier-delay** command.

```
Router#configure
Router(config)#interface HundredGigE 0/0/0/0
Router(config-if)#carrier-delay up 1000 down 150
Router(config-if)#commit
```

Running Configuration

```
interface HundredGigE0/0/0/0
carrier-delay up 1000 down 150
!
```

Verification

Run the **show interfaces** command to see the current state of the carrier-delay configuration for an interface.

```
Router#show interfaces HundredGigE 0/0/0/0 | include Carrier
Fri Mar 31 07:25:05.273 UTC
Carrier delay (up) is 1000 msec, Carrier delay (down) is 150 msec
```

How to Configure Ethernet

This section provides the following configuration procedures:

Configuring LLDP



Note LLDP is not supported on the FP-X line cards.

This section includes the following configuration topics for LLDP:

LLDP Default Configuration

This table shows the values of the LLDP default configuration on the router. To change the default settings, use the LLDP global configuration and LLDP interface configuration commands.

LLDP Function	Default
LLDP global state	Disabled
LLDP holdtime (before discarding)	120 seconds
LLDP timer (packet update frequency)	30 seconds
LLDP reinitialization delay	2 seconds
LLDP TLV selection	All TLVs are enabled for sending and receiving.
LLDP interface state	Enabled for both transmit and receive operation when LLDP is globally enabled.

Enabling LLDP Per Interface

When you enable LLDP globally, all interfaces that support LLDP are automatically enabled for both transmit and receive operations. However, if you want to enable LLDP per interface, perform the following configuration steps:

```
RP/0/RSP0/CPU0:ios(config)# int HundredGigE 0/2/0/0
RP/0/RSP0/CPU0:ios(config-if)# no sh
RP/0/RSP0/CPU0:ios(config-if)#commit
RP/0/RSP0/CPU0:ios(config-if)#lldp ?
RP/0/RSP0/CPU0:ios(config-if)#lldp enable
RP/0/RSP0/CPU0:ios(config-if)#commit
```

Running configuration

```
RP/0/RSP0/CPU0:ios#sh running-config
Wed Jun 27 12:40:21.274 IST
Building configuration...
!! IOS XR Configuration 0.0.0
!! Last configuration change at Wed Jun 27 00:59:29 2018 by UNKNOWN
!
interface HundredGigE0/1/0/0
 shutdown
!
interface HundredGigE0/1/0/1
 shutdown
!
interface HundredGigE0/1/0/2
 shutdown
!
```

```

interface HundredGigE0/2/0/0
 Shutdown
!
interface HundredGigE0/2/0/1
 shutdown
!
interface HundredGigE0/2/0/2
 shutdown
!
end

```

Verification

Verifying the config

```

=====
RP/0/RSP0/CPU0:ios#sh lldp interface <===== LLDP enabled only on GigEth0/2/0/0
Wed Jun 27 12:43:26.252 IST

```

```

HundredGigE0/2/0/0:
  Tx: enabled
  Rx: enabled
  Tx state: IDLE
  Rx state: WAIT FOR FRAME
RP/0/RSP0/CPU0:ios#

```

```

RP/0/RSP0/CPU0:ios# show lldp neighbors
Wed Jun 27 12:44:38.977 IST

```

Capability codes:

```

(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

```

```

Device ID      Local Intf      Hold-time  Capability  Port ID
ios            Gi0/2/0/0      120        R           Gi0/2/0/0    <===== LLDP
enabled only on GigEth0/2/0/0 and neighborhood seen for the same.

```

Total entries displayed: 1

```

RP/0/RSP0/CPU0:ios#

```

Enabling LLDP Globally

To run LLDP on the router, you must enable it globally. When you enable LLDP globally, all interfaces that support LLDP are automatically enabled for both transmit and receive operations.

You can override this default operation at the interface to disable receive or transmit operations. For more information about how to selectively disable LLDP receive or transmit operations for an interface, see the *Disabling LLDP Receive and Transmit Operation for an Interface* section.



Note For LLDP to work on any bundle member, enable LLDP on the bundle main interface either globally or on the interface itself. You can then choose to disable LLDP transmission on bundle main interface by using the `lldp transmit disable` command.

You can also control LLDP transmit or receive on each bundle member interface as desired.

The following table describes the global attributes that you can configure:

Attribute	Default	Range	Description
Holdtime	120	0-65535	Specifies the holdtime (in sec) that are sent in packets
Reinit	2	2-5	Delay (in sec) for LLDP initialization on any interface
Timer	30	5-65534	Specifies the rate at which LLDP packets are sent (in sec)

To enable LLDP globally, complete the following steps:

1. RP/0//CPU0:router # configure
2. RP/0//CPU0:router(config) #lldp
3. end or commit

Running configuration

```
RP/0/RP0/CPU0:turin-5#show run lldp
Fri Dec 15 20:36:49.132 UTC
lldp
!
```

```
RP/0/RP0/CPU0:turin-5#show lldp neighbors
Fri Dec 15 20:29:53.763 UTC
Capability codes:
```

```
(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
```

```
Device ID      Local Intf          Hold-time  Capability  Port ID
SW-NOSTG-I11-PUB.cis Mg0/RP0/CPU0/0    120       N/A        Fa0/28
```

Total entries displayed: 1

```
RP/0/RP0/CPU0:turin-5#show lldp neighbors mgmtEth 0/RP0/CPU0/0
Fri Dec 15 20:30:54.736 UTC
```

Capability codes:

```
(R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other
```

```
Device ID      Local Intf          Hold-time  Capability  Port ID
SW-NOSTG-I11-PUB.cis Mg0/RP0/CPU0/0    120       N/A        Fa0/28
```

Total entries displayed: 1

Configuring Global LLDP Operational Characteristics

When you enable LLDP globally on the router using the **lldp** command, these defaults are used for the protocol.

To modify the global LLDP operational characteristics such as the LLDP neighbor information holdtime, initialization delay, or packet rate, complete the following steps:

Procedure

Step 1 Example:

```
/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **lldp holdtime** *seconds*

Example:

```
RP/0//CPU0:router(config)#lldp holdtime 60
```

(Optional) Specifies the length of time that information from an LLDP packet should be held by the receiving device before aging and removing it.

Step 3 **lldp reinit** *seconds*

Example:

```
RP/0//CPU0:router(config)# lldp reinit 4
```

(Optional) Specifies the length of time to delay initialization of LLDP on an interface.

Step 4 **lldp timer** *seconds*

Example:

```
RP/0//CPU0:router(config)#lldp timer 60
```

(Optional) Specifies the LLDP packet rate.

Step 5 **end** or **commit**

Example:

```
RP/0//CPU0:router(config)# end
```

or

```
RP/0//CPU0:router(config)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?  
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.

- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Disabling Transmission of Optional LLDP TLVs

Certain TLVs are classified as mandatory in LLDP packets, such as the Chassis ID, Port ID, and Time to Live (TTL) TLVs. These TLVs must be present in every LLDP packet. You can suppress transmission of certain other optional TLVs in LLDP packets.

To disable transmission of optional LLDP TLVs, complete the following steps:

Procedure

Step 1 **configure**

Example:

```
RP/0/RSP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **lldp tlv-select *tlv-name* disable**

Example:

```
RP/0/RSP0/CPU0:router(config)# lldp tlv-select system-capabilities disable
```

(Optional) Specifies that transmission of the selected TLV in LLDP packets is disabled. The *tlv-name* can be one of the following LLDP TLV types:

- **management-address**
- **port-description**
- **system-capabilities**
- **system-description**
- **system-name**

Step 3 **end or commit**

Example:

```
RP/0/RSP0/CPU0:router(config)# end
```

or

```
RP/0/RSP0RSP0/CPU0:router(config)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?  
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Disabling LLDP Receive and Transmit Operation for an Interface

When you enable LLDP globally on the router, all supported interfaces are automatically enabled for LLDP receive and transmit operation. You can override this default by disabling these operations for a particular interface.

To disable LLDP receive and transmit operations for an interface, complete the following steps:

Procedure

Step 1 **configure**

Example:

```
RP/0/RSP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface HundredGigE 0/2/0/0**

Example:

```
RP/0/RSP0RP0/CPU0:router (config) #interface HundredGigE 0/2/0/0
```

Enters interface configuration mode and specifies the Ethernet interface name and notation *rack/slot/module/port*. Possible interface types for this procedure are:

- HundredGigE
- TenGigE

Step 3 **lldp**

Example:

```
RP/0/RSP0/CPU0:router(config-if)#lldp
```

(Optional) Enters LLDP configuration mode for the specified interface.

Step 4 receive disable

Example:

```
RP/0/RSP0/CPU0:router(config-lldp)#receive disable
```

(Optional) Disables LLDP receive operations on the interface.

Step 5 transmit disable

Example:

```
RP/0/RSP0/CPU0:router(config-lldp)#transmit disable
```

(Optional) Disables LLDP transmit operations on the interface.

Step 6 end or commit

Example:

```
RP/0/RSP0/CPU0:router(config)# end
```

or

```
RP/0/RSP0/CPU0:router(config)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Verifying the LLDP Configuration

This section describes how you can verify the LLDP configuration both globally and for a particular interface.

Verifying the LLDP Global Configuration

To verify the LLDP global configuration status and operational characteristics, use the **show lldp** command as shown in the following example:

```
RP/0/RSP0/CPU0:router# show lldp
Wed Apr 13 06:16:45.510 DST
Global LLDP information:
  Status: ACTIVE
  LLDP advertisements are sent every 30 seconds
  LLDP hold time advertised is 120 seconds
  LLDP interface reinitialisation delay is 2 seconds
```

If LLDP is not enabled globally, the following output appears when you run the **show lldp** command:

```
RP/0/RSP0/CPU0:router# show lldp
Wed Apr 13 06:42:48.221 DST
% LLDP is not enabled
```

Verifying the LLDP Interface Configuration

To verify the LLDP interface status and configuration, use the **show lldp interface** command as shown in the following example:

```
RP/0/RSP0/CPU0:router# show lldp interface HundredGigE 0/1/0/7
Wed Apr 13 13:22:30.501 DST

HundredGigE0/1/0/7:
  Tx: enabled
  Rx: enabled
  Tx state: IDLE
  Rx state: WAIT FOR FRAME
```

To monitor and maintain LLDP on the system or get information about LLDP neighbors, use one of the following commands:

Command	Description
clear lldp counters	Resets LLDP traffic counters or LLDP neighbor information.
show lldp entry	Displays detailed information about LLDP neighbors.
show lldp errors	Displays LLDP error and overflow statistics.
show lldp neighbors	Displays information about LLDP neighbors.
show lldp traffic	Displays statistics for LLDP traffic.

To collect or clear LLDP interface statistics, you can use the following commands:

Command	Description
show lldp traffic interface <i>interface_name</i>	Displays LLDP traffic statistics for the specified interface.

Command	Description
<code>clear lldp counters interface <i>interface_name</i></code>	Clears LLDP traffic statistics for the specified interface. Global statistics remains intact. Similarly, clearing global statistics does not impact the interface statistics.

Examples for LLDP Interface Statistics

This example shows interface statistics for **gigabitEthernet0/0/0/0**:

```
Router#show lldp traffic interface gigabitEthernet0/0/0/0
```

This example clears the interface statistics for **gigabitEthernet0/0/0/0**.

```
Router#show lldp traffic interface gigabitEthernet0/0/0/0
```

Running Configuration

```
Router#show lldp traffic interface gigabitEthernet 0/2/0/8
Wed Aug 24 17:38:11.829 IST
```

```
LLDP Interface statistics:
  Total frames out: 28786
  Total frames in: 38417
  Total frames received in error: 0
  Total frames out error: 0
  Total frames discarded: 0
  Total TLVs discarded: 0
  Total TLVs unrecognized: 0
```

Configuring LLDP Snoop

If you have LLDP enabled on all Ethernet interfaces, the system enables Link Layer Discovery Protocol (LLDP) snoop on all L2 interfaces by default. You can use LLDP snooping to troubleshoot problems at the client ports.



Note LLDP snoop is enabled only when LLDP RX is enabled and LLDP TX (transmit) is disabled either on interface or global LLDP configuration.

To enable LLDP snoop on an L2 interface, perform the following steps:

```
RP/0/RSP0/CPU0:ios# configure
RP/0/RSP0/CPU0:ios(config)# interface FourHundredGigE 0/0/0/5
RP/0/RSP0/CPU0:ios(config-if)#lldp
RP/0/RSP0/CPU0:ios(config-if)#enable
RP/0/RSP0/CPU0:ios(config-if)#transmit disable
RP/0/RSP0/CPU0:ios(config-if)#commit
```

Running Configuration

```

RP/0/RP0/CPU0:router#show run
Fri Jan 21 17:45:17.529 UTC
Building configuration...
!! IOS XR Configuration 7.7.1.06I
!! Last configuration change at Fri Jan 21 17:20:27 2022 by cisco
!
hostname router1
logging console disable
username xxxx
  group root-lr
  group cisco-support
  secret 10
$6$JELNK0oJaZZN7K0.$8YmyRWkq3D92i.lJc5QsdKq4kUjU.g9U7sYIIAV1QVnSBemng5q.5EyYv6xSL9niDxRmKaFEATs9BkitDqpr.
!
line console
  exec-timeout 0 0
  absolute-timeout 0
  session-timeout 0
!
line default
  exec-timeout 0 0
  absolute-timeout 0
  session-timeout 0
!
vty-pool default 0 99 line-template default
call-home
service active
contact smart-licensing
profile CiscoTAC-1
  active
  destination transport-method email disable
  destination transport-method http
!
!
interface MgmtEth0/RP0/CPU0/0
  shutdown
!
interface FourHundredGigE0/0/0/0
  lldp
    enable
    transmit disable
  !
  l2transport
  !
!
interface FourHundredGigE0/0/0/1
  shutdown
!
interface FourHundredGigE0/0/0/2
  shutdown
!
interface FourHundredGigE0/0/0/3
  shutdown
!
interface FourHundredGigE0/0/0/4
  shutdown
!
interface FourHundredGigE0/0/0/5
  lldp
    enable
    transmit disable
  !

```

```
l2transport
!  
!  
interface FourHundredGigE0/0/0/6
 shutdown
!  
interface FourHundredGigE0/0/0/7
 shutdown
!  
interface FourHundredGigE0/0/0/8
 shutdown
!  
interface FourHundredGigE0/0/0/9
 shutdown
!  
interface FourHundredGigE0/0/0/10
 shutdown
!  
interface FourHundredGigE0/0/0/11
 shutdown
!  
interface FourHundredGigE0/0/0/12
 shutdown
!  
interface FourHundredGigE0/0/0/13
 shutdown
!  
interface FourHundredGigE0/0/0/14
 shutdown
!  
interface FourHundredGigE0/0/0/15
 shutdown
!  
interface FourHundredGigE0/0/0/16
 shutdown
!  
interface FourHundredGigE0/0/0/17
 shutdown
!  
interface FourHundredGigE0/0/0/18
 shutdown
!  
interface FourHundredGigE0/0/0/19
 shutdown
!  
interface FourHundredGigE0/0/0/20
 shutdown
!  
interface FourHundredGigE0/0/0/21
 shutdown
!  
interface FourHundredGigE0/0/0/22
 shutdown
!  
interface FourHundredGigE0/0/0/23
 shutdown
!  
interface HundredGigE0/0/0/24
 shutdown
!  
interface HundredGigE0/0/0/25
 shutdown
!  
interface HundredGigE0/0/0/26
```



```

    shutdown
    !
interface HundredGigE0/0/0/27
    shutdown
    !
interface HundredGigE0/0/0/28
    shutdown
    !
interface HundredGigE0/0/0/29
    shutdown
    !
interface HundredGigE0/0/0/30
    shutdown
    !
interface HundredGigE0/0/0/31
    shutdown
    !
interface HundredGigE0/0/0/32
    shutdown
    !
interface HundredGigE0/0/0/33
    shutdown
    !
interface HundredGigE0/0/0/34
    shutdown
    !
interface HundredGigE0/0/0/35
    shutdown
    !
l2vpn
    bridge group bg1
        bridge-domain bd1
            interface FourHundredGigE0/0/0/0
                !
            interface FourHundredGigE0/0/0/5
                !
            !
        !
    !
end

RP/0/RP0/CPU0:router#

```

Verification

```

router0 <---> router1 <---> router2
         0/0/0/0         0/0/0/0/5

```

```

RP/0/RP0/CPU0:router0#config
Fri Jan 21 17:16:41.713 UTC
RP/0/RP0/CPU0:router0(config)#lldp
RP/0/RP0/CPU0:router0(config-lldp)#exit
RP/0/RP0/CPU0:router0(config)#int hu 0/0/0/0
RP/0/RP0/CPU0:router0(config-if)#no shut
RP/0/RP0/CPU0:router0(config-if)#end
Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:yes

```

```

RP/0/RP0/CPU0:router1#config
Fri Jan 21 17:17:41.459 UTC
RP/0/RP0/CPU0:router1(config)#int FourHundredGigE 0/0/0/0
RP/0/RP0/CPU0:router1(config-if)#no shut
RP/0/RP0/CPU0:router1(config-if)#l2transport
RP/0/RP0/CPU0:router1(config-if-l2)#exit
RP/0/RP0/CPU0:router1(config-if)#lldp

```

```

RP/0/RP0/CPU0:router1(config-lldp)#enable
RP/0/RP0/CPU0:router1(config-lldp)#transmit disable
RP/0/RP0/CPU0:router1(config-lldp)#exit
RP/0/RP0/CPU0:router1(config-if)#exit
RP/0/RP0/CPU0:router1(config)#int FourHundredGigE 0/0/0/5
RP/0/RP0/CPU0:router1(config-if)#no shut
RP/0/RP0/CPU0:router1(config-if)#l2transport
RP/0/RP0/CPU0:router1(config-if-l2)#exit
RP/0/RP0/CPU0:router1(config-if)#lldp
RP/0/RP0/CPU0:router1(config-lldp)#enable
RP/0/RP0/CPU0:router1(config-lldp)#transmit disable
RP/0/RP0/CPU0:router1(config-lldp)#exit
RP/0/RP0/CPU0:router1(config-if)#exit
RP/0/RP0/CPU0:router1(config)#l2vpn bridge group bg1
RP/0/RP0/CPU0:router1(config-l2vpn-bg)#bridge-domain bd1
RP/0/RP0/CPU0:router1(config-l2vpn-bg-bd)#interface FourHundredGigE 0/0/0/0
RP/0/RP0/CPU0:router1(config-l2vpn-bg-bd-ac)#exit
RP/0/RP0/CPU0:router1(config-l2vpn-bg-bd)#interface FourHundredGigE 0/0/0/5
RP/0/RP0/CPU0:router1(config-l2vpn-bg-bd-ac)#end
Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:yes

```

```

RP/0/RP0/CPU0:router0#config
Fri Jan 21 17:16:41.713 UTC
RP/0/RP0/CPU0:router0(config)#lldp
RP/0/RP0/CPU0:router0(config-lldp)#exit
RP/0/RP0/CPU0:router0(config)#int hu 0/0/0/0
RP/0/RP0/CPU0:router0(config-if)#no shut
RP/0/RP0/CPU0:router0(config-if)#end
Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:yes

```

```

RP/0/RP0/CPU0:router0#sh lldp neighbors
Fri Jan 21 17:21:15.857 UTC
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID          Local Intf                Hold-time  Capability  Port ID
router2            HundredGigE0/0/0/0       120        R
FourHundredGigE0/0/0/5

Total entries displayed: 1

```

```

RP/0/RP0/CPU0:router0#

RP/0/RP0/CPU0:router0#sh lldp neighbors
Fri Jan 21 17:21:15.857 UTC
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
  (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID          Local Intf                Hold-time  Capability  Port ID
router2            HundredGigE0/0/0/0       120        R
FourHundredGigE0/0/0/5

Total entries displayed: 1

```

```

RP/0/RP0/CPU0:router0#

RP/0/RP0/CPU0:router2#sh lldp neighbors
Fri Jan 21 17:21:20.998 UTC
Capability codes:
  (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device

```

```

(W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

Device ID          Local Intf          Hold-time  Capability  Port ID
router0            FourHundredGigE0/0/0/5  120       R
HundredGigE0/0/0/0

Total entries displayed: 1

RP/0/RP0/CPU0:router2#

RP/0/RP0/CPU0:router1#show controllers npu stats traps-all instance all location all | inc
LLDP
Fri Jan 21 17:24:07.964 UTC
LLDP
    3975      IFG      1520      0          0  22  RPLC_CPU  206  1538  6  4000
LLDP_SNOOP
    3862      NPU      N/A       16         0  28  RPLC_CPU  206  1538  6  4000
RP/0/RP0/CPU0:router1#

```

Configuration Examples for Ethernet

This section provides the following configuration examples:

Configuring an Ethernet Interface: Example

The following example shows how to configure an interface for a 10-Gigabit Ethernet modular services card:

```

RP/0//CPU0:router# configure
RP/0//CPU0:router(config)# interface TenGigE 0/0/0/1
RP/0//CPU0:router(config-if)# ipv4 address 172.18.189.38 255.255.255.224
RP/0//CPU0:router(config-if)# flow-control ingress
RP/0//CPU0:router(config-if)# mtu 1448
RP/0//CPU0:router(config-if)# mac-address 0001.2468.ABCD
RP/0//CPU0:router(config-if)# no shutdown
RP/0//CPU0:router(config-if)# end
Uncommitted changes found, commit them? [yes]: yes

RP/0//CPU0:router# show interfaces TenGigE 0/0/0/1

TenGigE0/0/0/1 is down, line protocol is down
  Hardware is TenGigE, address is 0001.2468.abcd (bia 0001.81a1.6b23)
  Internet address is 172.18.189.38/27
  MTU 1448 bytes, BW 10000000 Kbit
    reliability 0/255, txload Unknown, rxload Unknown
  Encapsulation ARPA,
    Full-duplex, 10000Mb/s, L/R
    output flow control is on, input flow control is on
  Encapsulation ARPA,
  ARP type ARPA, ARP timeout 01:00:00
  Last clearing of "show interface" counters never
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
    0 packets input, 0 bytes, 0 total input drops
    0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 0 multicast packets
    0 runts, 0 giants, 0 throttles, 0 parity

```

```

0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
0 packets output, 0 bytes, 0 total output drops
Output 0 broadcast packets, 0 multicast packets
0 output errors, 0 underruns, 0 applique, 0 resets
0 output buffer failures, 0 output buffers swapped out
0 carrier transitions

```

Configuring LLDP: Examples

The following example shows how to enable LLDP globally on the router and modify the default LLDP operational characteristics:

```

RP/0//CPU0:router# configure
RP/0//CPU0:router(config)# lldp
RP/0//CPU0:router(config)# lldp holdtime 60
RP/0//CPU0:router(config)# lldp reinit 4
RP/0//CPU0:router(config)# lldp timer 60
RP/0//CPU0:router(config)# commit

```

The following example shows how to disable a specific Gigabit Ethernet interface for LLDP transmission:

```

RP/0//CPU0:router# configure
RP/0//CPU0:router(config)# interface HundredGigE 0/2/0/0
RP/0//CPU0:router(config-if)# lldp
RP/0//CPU0:router(config-lldp)# transmit disable

```

Where to Go Next

When you have configured an Ethernet interface, you can configure individual VLAN subinterfaces on that Ethernet interface.

For information about modifying Ethernet management interfaces for the shelf controller (SC), route processor (RP), and distributed RP, see the Advanced Configuration and Modification of the Management Ethernet Interface later in this document.

For information about IPv6 see the Implementing Access Lists and Prefix Lists on Cisco IOS XR Software module in the Cisco IOS XR IP Addresses and Services Configuration Guide.

Configuring a Layer 2 VPN AC: Example

The following example indicates how to configure a Layer 2 VPN AC on an Ethernet interface:

```

RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# interface TenGigE 0/0/0/2
RP/0/RSP0/CPU0:router(config-if)# l2transport
RP/0/RSP0/CPU0:router(config-if-l2)# l2protocol tunnel
RP/0/RSP0/CPU0:router(config-if-l2)# commit

```

Configuring Physical Ethernet Interfaces

Use this procedure to create a basic Ethernet interface configuration.

Procedure

Step 1 **show version**

Example:

```
RP/0/RP0/CPU0:router# show version
```

(Optional) Displays the current software version, and can also be used to confirm that the router recognizes the line card.

Step 2 **show interfaces [TenGigE FortyGigE HundredGigE FourHundredGigE] interface-path-id**

Example:

```
RP/0/RP0/CPU0:router# show interface HundredGigE 0/1/0/1
```

(Optional) Displays the configured interface and checks the status of each interface port.

Step 3 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure terminal
```

Enters global configuration mode.

Step 4 **show interfaces [TenGigE FortyGigE HundredGigE FourHundredGigE] interface-path-id**

Example:

```
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/1
```

Enters interface configuration mode and specifies the Ethernet interface name and notation *rack/slot/module/port*. Possible interface types for this procedure are:

- 10GigE
- 40GigE
- 100GigE

Note

- The example indicates a 100-Gigabit Ethernet interface in the line card in slot 1.

- 400GigE

The examples of *interface-path-id* ranges are:

- **TenGigE** — 0/0/0/0 - 0/0/0/31
- **FortyGigE** — 0/0/1/0 - 0/0/1/1
- **HundredGigE** — 0/0/1/0 - 0/0/1/1

Step 5 **ipv4 address ip-address mask**

Example:

```
RP/0/RP0/CPU0:router(config-if)# ipv4 address 172.18.189.38 255.255.255.224
```

Assigns an IP address and subnet mask to the interface.

- Replace *ip-address* with the primary IPv4 address for the interface.
- Replace *mask* with the mask for the associated IP subnet. The network mask can be specified in either of two ways:
 - The network mask can be a four-part dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means that the corresponding address bit belongs to the network address.
 - The network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are network address.

Step 6 **flow-control {bidirectional| egress | ingress}****Example:**

```
RP/0/RP0/CPU0:router(config-if)# flow control ingress
```

(Optional) Enables the sending and processing of flow control pause frames.

- **egress**—Enables the sending of flow control pause frames in egress.
- **ingress**—Enables the processing of received pause frames on ingress.
- **bidirectional**—Enables the sending of flow control pause frames in egress and the processing of received pause frames on ingress.

Step 7 **mtu bytes****Example:**

```
RP/0/RP0/CPU0:router(config-if)# mtu 1448
```

(Optional) Sets the MTU value for the interface.

- The default is 1514 bytes for normal frames and 1518 bytes for 802.1Q tagged frames.
- The range for 100-Gigabit Ethernet mtu values is 64 bytes to 65535 bytes.

Step 8 **no shutdown****Example:**

```
RP/0/RP0/CPU0:router(config-if)# no shutdown
```

Removes the shutdown configuration, which forces an interface administratively down.

Step 9 **end or commit****Example:**

```
RP/0/RP0/CPU0:router(config-if)# end
```

or

```
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Step 10 **show interfaces [TenGigE FortyGigE HundredGigE FourHundredGigE] interface-path-id**

Example:

```
RP/0/RP0/CPU0:router# show interfaces HundredGigE 0/1/0/1
```

(Optional) Displays statistics for interfaces on the router.

Example

This example shows how to configure an interface for a 100-Gigabit Ethernet line card:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/1
RP/0/RP0/CPU0:router(config-if)# ipv4 address 172.18.189.38 255.255.255.224

RP/0/RP0/CPU0:router(config-if)# mtu 1448

RP/0/RP0/CPU0:router(config-if)# no shutdown
RP/0/RP0/CPU0:router(config-if)# end
Uncommitted changes found, commit them? [yes]: yes

RP/0/RP0/CPU0:router# show interfaces HundredGigE 0/5/0/24
HundredGigE0/5/0/24 is up, line protocol is up
  Interface state transitions: 1
  Hardware is HundredGigE, address is 6219.8864.e330 (bia 6219.8864.e330)
  Internet address is 3.24.1.1/24
  MTU 9216 bytes, BW 100000000 Kbit (Max: 100000000 Kbit)
    reliability 255/255, txload 3/255, rxload 3/255
  Encapsulation ARPA,
  Full-duplex, 100000Mb/s, link type is force-up
  output flow control is off, input flow control is off
```

```

Carrier delay (up) is 10 msec
loopback not set,
Last link flapped 10:05:07
ARP type ARPA, ARP timeout 04:00:00
Last input 00:08:56, output 00:00:00
Last clearing of "show interface" counters never
5 minute input rate 1258567000 bits/sec, 1484160 packets/sec
5 minute output rate 1258584000 bits/sec, 1484160 packets/sec
228290765840 packets input, 27293508436038 bytes, 0 total input drops
0 drops for unrecognized upper-level protocol
Received 15 broadcast packets, 45 multicast packets
0 runts, 0 giants, 0 throttles, 0 parity
0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort
212467849449 packets output, 25733664696650 bytes, 0 total output drops
Output 23 broadcast packets, 15732 multicast packets
39 output errors, 0 underruns, 0 applique, 0 resets
0 output buffer failures, 0 output buffers swapped out
0 carrier transitions

```

```
RP/0/RP0/CPU0:router# show running-config interface HundredGigE 0/5/0/24
```

```

interface HundredGigE 0/5/0/24
mtu 9216
service-policy input linerate
service-policy output elinerate
ipv4 address 3.24.1.1 255.255.255.0
ipv6 address 3:24:1::1/64
flow ipv4 monitor perfv4 sampler fsm ingress
!

```

Viewing Interface Counters Report

The Interface Counters report summarizes the statistics for all interfaces configured on the router.

The report displays the interfaces configured, the input and output rate, the total number of packets transmitted and received, the time interval, the current status of each interface, and the packet counts for input and output broadcast, multicast, and errored packets.

The **show interfaces** command, displays statistics per interface with many lines of data. The traffic rate displays the average number of packets received per second over the load interval. The load interval is configurable on the physical and bundle main interface. The report displays the load on the interface for a longer duration of time and does not show a spike in the traffic rate. This rate is the exponentially weighted average with a time constant of the load interval.



Note For the average to be within two percent of the instantaneous rate of a uniform stream of traffic, four times the load interval must pass.

For more information about the use of **show interfaces** command, see *Interface and Hardware Component Command Reference for Cisco 8000 Series Routers*.

Instant Display of Traffic Rates for all the Physical Interfaces

Table 13: Feature History Table

Feature Name	Release Information	Feature Description
Instant display of traffic rates on all physical interfaces	Release 7.5.4	<p>You can now display a snapshot of the traffic throughput and traffic rate on all physical interfaces over the last few seconds. We have introduced a show command to view the counters and rate information for the interfaces.</p> <p>The feature introduces these:</p> <ul style="list-style-type: none"> • CLI: show interfaces counter rates physical • YANG Data Model: Cisco-IOS-XR-infra-statsd-oper.yang (see GitHub under the 754 folder.)

The new **show** command displays a snapshot of statistics for all the interfaces at a given instant for your quick reference. Here, the display is in a tabular format for easy analysis.

Run the **show interfaces counter rates physical** command to view statistics of all physical interfaces.

View the statistics

```
Router#show interfaces counters rates physical
```

InterfaceName	Intvl	InMbps	InBW%	InKpps	OutMbps	OutBW%	OutKpps
GigabitEthernet0/2/0/0	0:05	0.0	0.0%	0.0	0.0	0.0%	0.0
GigabitEthernet0/2/0/1	0:05	0.0	0.0%	0.0	0.0	0.0%	0.0
GigabitEthernet0/2/0/2	0:05	0.0	0.0%	0.0	0.0	0.0%	0.0
GigabitEthernet0/2/0/3	0:05	235.0	22.0%	23.5	87.0	9.5%	7.2
GigabitEthernet0/3/0/0	0:05	88.0	9.3%	7.0	100.0	10.0%	10.5
GigabitEthernet0/3/0/1	0:05	0.0	0.0%	0.0	0.0	0.0%	0.0

The statistics for each physical interface is calculated for the time interval of 5 sec. Hence, the input and output rate (in Mbps and Kpps) is the real-time statistics.



Note The traffic rate displayed is the real-time link utilization of the time interval. The time interval is determined by the system and may vary based on the system processing load. The time interval increases during events where the system is handling, for example, performing routing updates.

How to Configure Interfaces in Breakout Mode

Information About Breakout

The router supports transmission of traffic in the breakout mode. The breakout mode enables a 40 Gigabit Ethernet port to be split into four independent and logical 10 Gigabit Ethernet ports. The 4x10 breakout mode is supported on the following types of 40G modules:

- QSFP-4x10-LR-S
- QSFP-40G-SR4

Guidelines and Restrictions for Breakout Mode

- The native 40G mode on QSFP-40G-SR4 is not supported.
- The 36-port QSFP56-DD 400 GbE Line Card does not support the 4x10G breakout.
- If you're using a Q100-based Cisco 8200 Series Router and want to set up a 4x10G breakout configuration, you need to use even numbered ports from 24 to 35. These include ports 24, 26, 28, 30, 32, and 34. Once you do this, the system automatically disables the odd numbered ports in this range - ports 25, 27, 29, 31, 33, and 35.
- Use the *hw-module port-range* command to set the port range for the breakout configuration in the global configuration.
- To remove the global *hw-module port-range* configuration, you must first remove the 'breakout 4x10' configuration under the controller.
- For 4x10G breakout on 48-port Line Card, only QSFP-4x10-LR-S module is supported.
- For the 88-LC0-34H14FH line card, you must breakout only 3 ports instead of 4 ports to avoid the QOS-DPA_QOSEA-2-TMPORT_PROG_ERROR issue, which creates partial interfaces during configuration mode.

Configure Breakout in a Port

Configuring breakout in a port:

```
RP/0/RP0/CPU0:uut# configure
Fri Oct 11 23:58:47.165 UTC
RP/0/RP0/CPU0:uut(config)# controller optics 0/1/0/28
RP/0/RP0/CPU0:uut(config-Optics)# breakout 4x10
RP/0/RP0/CPU0:uut(config-Optics)# commit
Fri Oct 11 23:59:51.261 UTC
RP/0/RP0/CPU0:uut(config-Optics)# end
RP/0/RP0/CPU0:uut#
```

Remove the Breakout Configuration

Removing the breakout configuration:

```
RP/0/RP0/CPU0:uut# configure
Sat Oct 12 00:01:38.673 UTC
RP/0/RP0/CPU0:uut(config)# controller optics 0/1/0/28
RP/0/RP0/CPU0:uut(config-Optics)# no breakout 4x10
RP/0/RP0/CPU0:uut(config-Optics)# commit
Sat Oct 12 00:01:55.864 UTC
RP/0/RP0/CPU0:uut(config-Optics)# end
```

Verify a Breakout Configuration

Verifying a breakout configuration:

```
RP/0/RP0/CPU0:uut# show running-config controller optics 0/1/0/28
Sat Oct 12 00:11:33.962 UTC
controller Optics0/1/0/28
breakout 4x10
!
```

```
RP/0/RP0/CPU0:uut# show int br location 0/1/CPU0 | i Te
Sat Oct 12 00:11:38.609 UTC
      Te0/1/0/27/0      up      up      ARPA 10000  10000000
      Te0/1/0/27/1      up      up      ARPA 10000  10000000
      Te0/1/0/27/2      up      up      ARPA 10000  10000000
      Te0/1/0/27/3      up      up      ARPA 10000  10000000
      Te0/1/0/28/0      up      up      ARPA 10000  10000000
      Te0/1/0/28/1      up      up      ARPA 10000  10000000
      Te0/1/0/28/2      up      up      ARPA 10000  10000000
      Te0/1/0/28/3      up      up      ARPA 10000  10000000
```

Ethernet Interface Route Statistics

Table 14: Feature History Table

Feature Name	Release Information	Feature Description
Ethernet Interface Route Statistics	Release 7.3.4	You can view statistics on the number of packets and bytes sent and received in unicast, multicast, and broadcast routes. These statistics help you to monitor the network performance and measure your bandwidth.

Ethernet interface route statistics provide the following information about the unicast, multicast, and broadcast routes:

- The number of packets or bytes received and transmitted.
- The total number of packets or bytes passing through the Ethernet interface.

These statistics are available on Cisco 8000 routers that are built with Cisco Silicon One Q100 and Q200 processors and all Network Processor Unit (NPU) 2.0 devices.

Ethernet interface route statistics may be useful for monitoring the network devices and their traffic. For example, if you are not able to connect to the internet or use some cloud-based applications, these route statistics can help you understand the problems in the network and where they occur.

Viewing the Interface Statistics

Use the `show interface` and the `show controller interface` commands to view these Ethernet interface route statistics. The following is a sample showing both commands.

```
Router#show interfaces HundredGigE 0/0/0/0
```

```
<Timestamp>
```

```
HundredGigE0/0/0/0 is up, line protocol is up
```

```
Interface state transitions: 93
```

```
Hardware is HundredGigE, address is acbc.d975.0500 (bia acbc.d975.0500)
```

```
Internet address is 100.0.1.1/24
```

```
MTU 1514 bytes, BW 100000000 Kbit (Max: 100000000 Kbit)
```

```
reliability 255/255, txload 0/255, rxload 0/255
```

```
Encapsulation ARPA,
```

```
Full-duplex, 100000Mb/s, 100GBASE-SR4, link type is force-up
```

```
output flow control is off, input flow control is off
```

```
Carrier delay (up) is 10 msec
```

```
loopback not set,
```

```
Last link flapped 01:03:01
```

```
ARP type ARPA, ARP timeout 04:00:00
```

```
Last input 3d09h, output 01:02:41
```

```
Last clearing of "show interface" counters never
```

```
5 minute input rate 0 bits/sec, 0 packets/sec
```

```
5 minute output rate 0 bits/sec, 0 packets/sec
```

```

2959131434 packets input, 757537646912 bytes, 0 total input drops

0 drops for unrecognized upper-level protocol

Received 0 broadcast packets, 0 multicast packets

    0 runts, 0 giants, 0 throttles, 0 parity

0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored, 0 abort

2958525230 packets output, 757382468319 bytes, 0 total output drops

Output 0 broadcast packets, 0 multicast packets

0 output errors, 0 underruns, 0 applique, 0 resets

0 output buffer failures, 0 output buffers swapped out

93 carrier transitions

```

Router#**show controllers HundredGigE0/0/0/0 stats**

<Timestamp>

Statistics for interface HundredGigE0/0/0/0 (cached values):

Ingress:

```

Input total bytes           = 757537646912

Input good bytes             = 757537646912

Input total packets        = 2959131434
Input 802.1Q frames         = 0
Input pause frames          = 0
Input pkts 64 bytes         = 1
Input pkts 65-127 bytes     = 0
Input pkts 128-255 bytes    = 0
Input pkts 256-511 bytes    = 2959131433
Input pkts 512-1023 bytes   = 0
Input pkts 1024-1518 bytes  = 0
Input pkts 1519-Max bytes   = 0

Input good pkts             = 2959131434

```

```

Input unicast pkts           = 0
Input multicast pkts        = 0
Input broadcast pkts       = 0

Input drop overrun           = 0
Input drop abort             = 0
Input drop invalid VLAN     = 0
Input drop invalid DMAC     = 0
Input drop invalid encap    = 0
Input drop other            = 0

Input error giant           = 0
Input error runt            = 0
Input error jabbers         = 0
Input error fragments       = 0
Input error CRC             = 0
Input error collisions      = 0
Input error symbol          = 0
Input error other           = 0

Input MIB giant             = 0
Input MIB jabber            = 0
Input MIB CRC               = 0

Egress:
Output total bytes         = 757382468319
Output good bytes           = 757382468319

Output total packets      = 2958525230
Output 802.1Q frames        = 0
Output pause frames         = 0
Output pkts 64 bytes        = 41
Output pkts 65-127 bytes    = 296
Output pkts 128-255 bytes   = 746
Output pkts 256-511 bytes   = 2958524147
Output pkts 512-1023 bytes  = 0
Output pkts 1024-1518 bytes = 0
Output pkts 1519-Max bytes  = 0

Output good pkts            = 2958525230
Output unicast pkts       = 0
Output multicast pkts    = 0
Output broadcast pkts    = 0

Output drop underrun        = 0
Output drop abort           = 0
Output drop other           = 0

Output error other          = 0

```



CHAPTER 6

Configuring Ethernet OAM

This module describes the configuration of Ethernet Operations, Administration, and Maintenance (OAM):

Table 15: Feature Information Table

Release	Modification
Release 7.3.1	Support for Ethernet Link OAM was introduced.

- [Information About Configuring Ethernet OAM, on page 69](#)
- [Configuration Examples for Ethernet OAM, on page 72](#)
- [Ethernet CFM, on page 74](#)
- [How to Configure Ethernet OAM, on page 86](#)
- [CFM Over Bundles, on page 108](#)
- [Ethernet Frame Delay Measurement for L2VPN Services, on page 109](#)
- [Link Loss Forwarding, on page 113](#)

Information About Configuring Ethernet OAM

To configure Ethernet OAM, you should understand the following concepts:

Ethernet Link OAM

Table 16: Feature History Table

Feature Name	Release Information	Feature Description
Ethernet Link OAM	Release 7.3.1	This feature allow Service Providers to monitor the quality of the connections on a MAN or WAN. Service providers can monitor specific events, and take actions on events. Ethernet link OAM operates on a single, physical link and it can be configured to monitor either side or both sides of that link.

Ethernet as a Metro Area Network (MAN) or a Wide Area Network (WAN) technology benefits greatly from the implementation of Operations, Administration and Maintenance (OAM) features. Ethernet link OAM features allow Service Providers to monitor the quality of the connections on a MAN or WAN. Service providers can monitor specific events, and take actions on events. Ethernet link OAM operates on a single, physical link and it can be configured to monitor either side or both sides of that link.

Ethernet link OAM can be configured in the following ways:

- A Link OAM profile can be configured, and this profile can be used to set the parameters for multiple interfaces.
- Link OAM can be configured directly on an interface.

When an interface is also using a link OAM profile, specific parameters that are set in the profile can be overridden by configuring a different value directly on the interface.

An Ethernet Link OAM profile simplifies the process of configuring EOAM features on multiple interfaces. An Ethernet OAM profile, and all of its features, can be referenced by other interfaces, allowing other interfaces to inherit the features of that Ethernet OAM profile.

Individual Ethernet link OAM features can be configured on individual interfaces without being part of a profile. In these cases, the individually configured features always override the features in the profile.

The preferred method of configuring custom EOAM settings is to create an EOAM profile in Ethernet configuration mode and then attach it to an individual interface or to multiple interfaces.

These standard Ethernet Link OAM features are supported on the router:

Neighbor Discovery

Neighbor discovery enables each end of a link to learn the OAM capabilities of the other end and establish an OAM peer relationship. Each end also can require that the peer have certain capabilities before it will establish a session. You can configure certain actions to be taken if there is a capabilities conflict or if a discovery process times out, using the **action capabilities-conflict** or **action discovery-timeout** commands.

EFD

Ethernet Fault Detection (EFD) is a mechanism that allows Ethernet OAM protocols to control the line protocol state of an interface.

Unlike many other interface types, Ethernet interfaces do not have a line protocol, whose state is independent from that of the interface. For Ethernet interfaces, this role is handled by the physical-layer Ethernet protocol itself, and therefore if the interface is physically up, then it is available and traffic can flow.

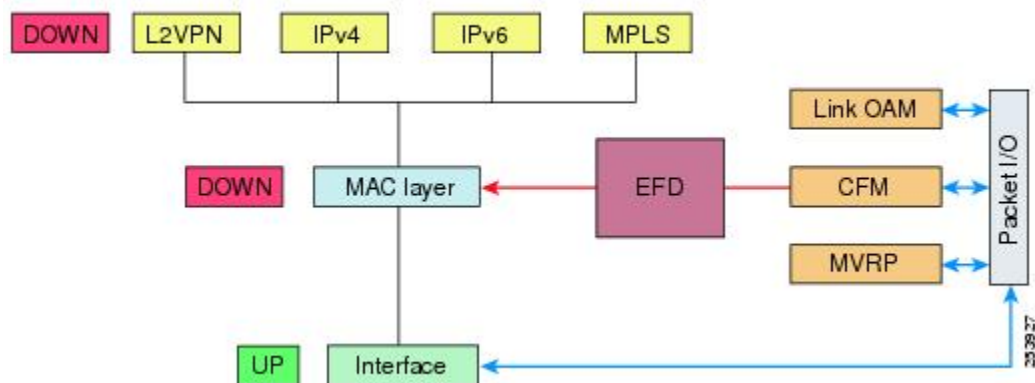
EFD changes this to allow EOAM to act as the line protocol for Ethernet interfaces. This allows EOAM to control the interface state so that if a EOAM defect (such as AIS or loss of continuity) is detected with an expected peer MEP, the interface can be shut down. This not only stops traffic flow, but also triggers actions in any higher-level protocols to route around the problem. For example, in the case of Layer 2 interfaces, the MAC table would be cleared and MSTP would reconverge. For Layer 3 interfaces, the ARP cache would be cleared and potentially the IGP would reconverge.



Note EFD can only be used for down MEPs. When EFD is used to shut down the interface, the EOAM frames continue to flow. This allows EOAM to detect when the problem has been resolved, and thus bring the interface backup automatically.

This figure shows EOAM detection of an error on one of its sessions EFD signaling an error to the corresponding MAC layer for the interface. This triggers the MAC to go to a down state, which further triggers all higher level protocols (Layer 2 pseudowires, IP protocols, and so on) to go down and also trigger a reconvergence where possible. As soon as EOAM detects there is no longer any error, it can signal to EFD and all protocols will once again go active.

Figure 1: EOAM Error Detection and EFD Trigger



MIB Retrieval

MIB retrieval enables an OAM peer on one side of an interface to get the MIB variables from the remote side of the link. The MIB variables that are retrieved from the remote OAM peer are READ ONLY.

Miswiring Detection (Cisco-Proprietary)

Miswiring Detection is a Cisco-proprietary feature that uses the 32-bit vendor field in every Information OAMPDU to identify potential miswiring cases.

SNMP Traps

SNMP traps can be enabled or disabled on an Ethernet OAM interface.

Configuration Examples for Ethernet OAM

This section provides the following configuration examples:

Configuring Ethernet OAM Features on an Individual Interface: Example

This example shows how to configure Ethernet OAM features on an individual interface:

```
configure
interface TenGigE 0/1/0/0
  ethernet oam
  link-monitor
  symbol-period window 60000
  symbol-period threshold ppm low 10000000 high 60000000
  frame window 60
  frame threshold ppm low 10000000 high 60000000
  frame-period window 60000
  frame-period threshold ppm low 100 high 12000000
  frame-seconds window 900000
  frame-seconds threshold low 3 high 900
  exit
mib-retrieval
connection timeout 30
require-remote mode active
require-remote mib-retrieval
action link-fault error-disable-interface
action dying-gasp error-disable-interface
action critical-event error-disable-interface
action discovery-timeout error-disable-interface
action session-down error-disable-interface
action capabilities-conflict error-disable-interface
action wiring-conflict error-disable-interface

commit
```

Configuring an Ethernet OAM Profile Globally: Example

This example shows how to configure an Ethernet OAM profile globally:

```
configure
ethernet oam profile Profile_1
  link-monitor
  symbol-period window 60000
  symbol-period threshold ppm low 10000000 high 60000000
  frame window 60
  frame threshold ppm low 10000000 high 60000000
```

```

frame-period window 60000
frame-period threshold ppm low 100 high 12000000
frame-seconds window 900000
frame-seconds threshold low 3 high 900
exit
mib-retrieval
connection timeout 30
require-remote mode active
require-remote mib-retrieval
action dying-gasp error-disable-interface
action critical-event error-disable-interface
action discovery-timeout error-disable-interface
action session-down error-disable-interface
action capabilities-conflict error-disable-interface
action wiring-conflict error-disable-interface

commit

```

Configuring Ethernet OAM Features to Override the Profile on an Individual Interface: Example

This example shows the configuration of Ethernet OAM features in a profile followed by an override of that configuration on an interface:

```

configure
ethernet oam profile Profile_1
mode passive
action dying-gasp disable
action critical-event disable
action discovery-timeout disable
action session-up disable
action session-down disable
action capabilities-conflict disable
action wiring-conflict disable

commit

configure
interface TenGigE 0/1/0/0
ethernet oam
profile Profile_1
mode active
action dying-gasp log
action critical-event log
action discovery-timeout log
action session-up log
action session-down log
action capabilities-conflict log
action wiring-conflict log

commit

```

Clearing Ethernet OAM Statistics on an Interface: Example

This example shows how to clear Ethernet OAM statistics on an interface:

```

RP/0/RP0/CPU0:router# clear ethernet oam statistics interface gigabitethernet 0/1/5/1

```

Enabling SNMP Server Traps on a Router: Example

This example shows how to enable SNMP server traps on a router:

```
configure
snmp-server traps ethernet oam events
```

Ethernet CFM

Table 17: Feature History Table

Feature Name	Release	Description
Monitoring Layer 3 Connectivity Using Down MEP on L3 Interfaces	Release 24.2.11	<p>This enhancement expands network diagnostics to L3 interfaces at L2 network termination, simplifying the management and maintenance of multilayer networks. Without impacting the underlying L2 infrastructure, this feature uses CFM packets to verify the connection of L3 paths.</p> <p>Previously, CFM Down MEP support was limited to L2 interfaces associated with cross-connect or bundle members.</p> <p>This feature is supported on both physical main and subinterfaces, bundle main and subinterfaces.</p>

Feature Name	Release	Description
CFM on Bundle Member Link for Connectivity Check	Release 7.3.15	<p>This feature introduces support for Connectivity Fault Management (CFM) on bundle members. Earlier, network administrators managed networks by using the fault, configuration, account, performance, security model. CFM is one of a suite of the Ethernet OAM protocols, which uses a combination of keepalive packets and MAC-based pings, and traceroutes to detect faults in a network.</p> <p>With the CFM feature, you:</p> <ul style="list-style-type: none"> • reduce operating expenses for service operators by reducing network faults and errors • provide end-to-end maintenance of networks
Up MEP and Down MEP Support in CFM	Release 7.3.15	<p>This feature introduces Maintenance End Points (MEP) entities that you can configure in a domain.</p> <p>MEPs send either CFM frames from the interface where they are configured or CFM frames that are received on other interfaces.</p> <p>MEPs allow you to perform fault management and carry out performance checks.</p>

Ethernet Connectivity Fault Management (CFM) is a service-level OAM protocol that provides tools for monitoring and troubleshooting end-to-end Ethernet services per VLAN. This includes proactive connectivity monitoring, fault verification, and fault isolation. CFM uses standard Ethernet frames and can be run on any physical media that is capable of transporting Ethernet service frames. Unlike most other Ethernet protocols which are restricted to a single physical link, CFM frames can transmit across the entire end-to-end Ethernet network.

CFM is defined in two standards:

- IEEE 802.1ag—Defines the core features of the CFM protocol.
- ITU-T Y.1731—Redefines, but maintains compatibility with the features of IEEE 802.1ag, and defines some additional features.

Ethernet CFM supports these functions of ITU-T Y.1731:

- ETH-CC, ETH-RDI, ETH-LB, ETH-LT—These are equivalent to the corresponding features defined in IEEE 802.1ag.



Note The Linktrace responder procedures defined in IEEE 802.1ag are used rather than the procedures defined in Y.1731; however, these are interoperable.

- ETH-AIS—The reception of ETH-LCK messages is also supported.

Limitations and Restrictions


- The system supports only cross-connect.
- MIPs are not supported.
- Supports timer of 1s, 10s, 1m, 10m.
- Supports timer of 100ms, 1s, 10s, 1m, 10m for bundle members.
- Multiple MEPs of different directions are not supported on the same interface or Xconnect.

Maintenance Domains

To understand how the CFM maintenance model works, you need to understand these concepts and features:

A maintenance domain describes a management space for the purpose of managing and administering a network. A domain is owned and operated by a single entity and defined by the set of interfaces internal to it and at its boundary, as shown in this figure.

Figure 2: CFM Maintenance Domain



A maintenance domain is defined by the bridge ports that are provisioned within it. Domains are assigned maintenance levels, in the range of 0 to 7, by the administrator. The level of the domain is useful in defining the hierarchical relationships of multiple domains.

CFM maintenance domains allow different organizations to use CFM in the same network, but independently. For example, consider a service provider who offers a service to a customer, and to provide that service, they use two other operators in segments of the network. In this environment, CFM can be used in the following ways:

- The customer can use CFM between their CE devices, to verify and manage connectivity across the whole network.
- The service provider can use CFM between their PE devices, to verify and manage the services they are providing.
- Each operator can use CFM within their operator network, to verify and manage connectivity within their network.

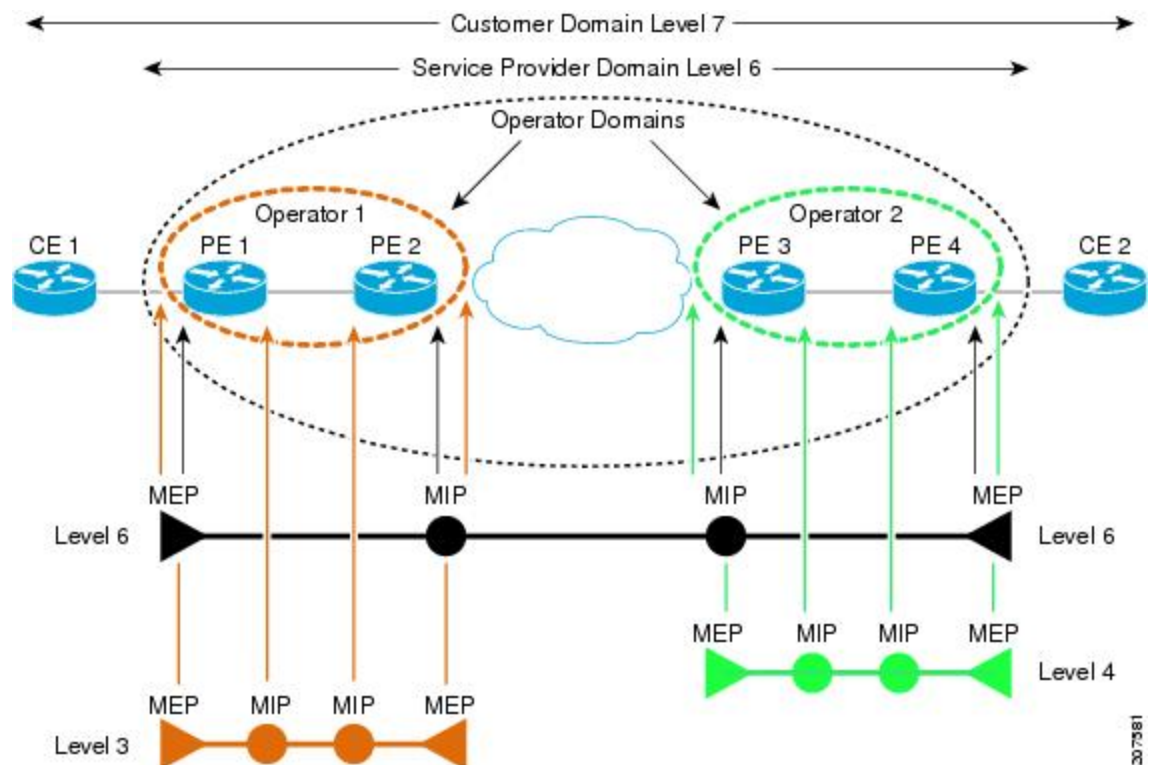
Each organization uses a different CFM maintenance domain.

This figure shows an example of the different levels of maintenance domains in a network.




Note In CFM diagrams, the conventions are that triangles represent MEPs, pointing in the direction that the MEP sends CFM frames, and circles represent MIPs.

Figure 3: Different CFM Maintenance Domains Across a Network



To ensure that the CFM frames for each domain do not interfere with each other, each domain is assigned a maintenance level, between 0 and 7. Where domains are nested, as in this example, the encompassing domain must have a higher level than the domain it encloses. In this case, the domain levels must be negotiated between the organizations involved. The maintenance level is carried in all CFM frames that relate to that domain.

CFM maintenance domains may touch or nest, but cannot intersect. This figure illustrates the supported structure for touching and nested domains, and the unsupported intersection of domains.



Services

A CFM service allows an organization to partition its CFM maintenance domain, according to the connectivity within the network. For example, if the network is divided into a number of virtual LANs (VLANs), a CFM service is created for each of these. CFM can then operate independently in each service. It is important that the CFM services match the network topology, so that CFM frames relating to one service cannot be received in a different service. For example, a service provider may use a separate CFM service for each of their customers, to verify and manage connectivity between that customer's end points.

A CFM service is always associated with the maintenance domain that it operates within, and therefore with that domain's maintenance level. All CFM frames relating to the service carry the maintenance level of the corresponding domain.



Note CFM Services are referred to as *Maintenance Associations* in IEEE 802.1ag and as *Maintenance Entity Groups* in ITU-T Y.1731.

Maintenance Points

A CFM Maintenance Point (MP) is an instance of a particular CFM service on a specific interface. CFM only operates on an interface if there is a CFM maintenance point on the interface; otherwise, CFM frames are forwarded transparently through the interface.

A maintenance point is always associated with a particular CFM service, and therefore with a particular maintenance domain at a particular level. Maintenance points generally only process CFM frames at the same level as their associated maintenance domain. Frames at a higher maintenance level are always forwarded transparently, while frames at a lower maintenance level are normally dropped. This helps enforce the maintenance domain hierarchy, and ensures that CFM frames for a particular domain cannot leak out beyond the boundary of the domain.

There are following type(s) of MP(s):

- **Maintenance End Points (MEPs)**—Created at the edge of the domain. Maintenance end points (MEPs) are members of a particular service within a domain and are responsible for sourcing and sinking CFM frames. They periodically transmit continuity check messages and receive similar messages from other MEPs within their domain. They also transmit traceroute and loopback messages at the request of the administrator. MEPs are responsible for confining CFM messages within the domain.

MEP and CFM Processing Overview

The boundary of a domain is an interface, rather than a bridge or host. Therefore, MEPs can be sub-divided into two categories:

- Down MEPs—Send CFM frames from the interface where they are configured, and process CFM frames received on that interface. Down MEPs transmit AIS messages upward (toward the cross-connect).
- Up MEPs—Send frames into the bridge relay function, as if they had been received on the interface where the MEP is configured. They process CFM frames that have been received on other interfaces, and have been switched through the bridge relay function as if they are going to be sent out of the interface where the MEP is configured. Up MEPs transmit AIS messages downward (toward the wire). However, AIS packets are only sent when there is a MIP configured on the same interface as the MEP and at the level of the MIP.

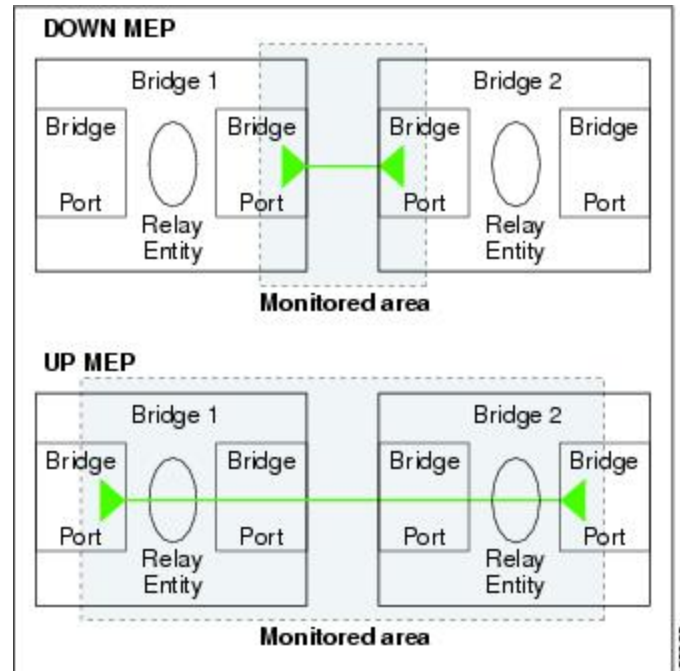


Note


- The terms *Down MEP* and *Up MEP* are defined in the IEEE 802.1ag and ITU-T Y.1731 standards, and refer to the direction that CFM frames are sent from the MEP. The terms should not be confused with the operational status of the MEP.
- The router only supports the “Down MEP level < Up MEP level” configuration.

This figure illustrates the monitored areas for Down and Up MEPs.

Figure 4: Monitored Areas for Down and Up MEPs



This figure shows maintenance points at different levels. Because domains are allowed to nest but not intersect, a MEP at a low level often corresponds with a MEP at a higher level.



Up MEPs can only exist on switched (Layer 2) interfaces, because they send and receive frames from the bridge relay function. Down MEPs can be created on switched (Layer 2) interfaces.

MEPs continue to operate normally if the interface they are created on is blocked by the Spanning Tree Protocol (STP); that is, CFM frames at the level of the MEP continue to be sent and received, according to the direction of the MEP. MEPs never allow CFM frames at the level of the MEP to be forwarded, so the STP block is maintained.



Note A separate set of CFM maintenance levels is created every time a VLAN tag is pushed onto the frame. Therefore, if CFM frames are received on an interface which pushes an additional tag, so as to “tunnel” the frames over part of the network, the CFM frames will not be processed by any MPs within the tunnel, even if they are at the same level. For example, if a CFM MP is created on an interface with an encapsulation that matches a single VLAN tag, any CFM frames that are received at the interface that have two VLAN tags will be forwarded transparently, regardless of the CFM level.

CFM Protocol Messages

The CFM protocol consists of a number of different message types, with different purposes. All CFM messages use the CFM EtherType, and carry the CFM maintenance level for the domain to which they apply.


This section describes the following CFM messages:

Continuity Check (IEEE 802.1ag and ITU-T Y.1731)

Continuity Check Messages (CCMs) are “heartbeat” messages exchanged periodically between all the MEPs in a service. Each MEP sends out multicast CCMs, and receives CCMs from all the other MEPs in the service—these are referred to as *peer MEPs*. This allows each MEP to discover its peer MEPs, and to verify that there is connectivity between them.

MIPs also receive CCMs. MIPs use the information to build a MAC learning database that is used when responding to Linktrace. For more information about Linktrace, see the [Linktrace \(IEEE 802.1ag and ITU-T Y.1731\)](#).

Figure 5: Continuity Check Message Flow



All the MEPs in a service must transmit CCMs at the same interval. IEEE 802.1ag defines the following possible intervals that can be used:

- 100 ms (only supported on bundle members)
- 1 s
- 10 s
- 1 minute
- 10 minutes

A MEP detects a loss of connectivity with one of its peer MEPs when some number of CCMs are missed. This occurs when sufficient time has passed during which a certain number of CCMs were expected, given the CCM interval. This number is called the *loss threshold*, and is usually set to 3.

With the exception of bundle members, CFM is supported only on interfaces that have Layer 2 transport feature enabled.

CCM messages carry a variety of information that allows different defects to be detected in the service. This information includes:

- A configured identifier for the domain of the transmitting MEP. This is referred to as the Maintenance Domain Identifier (MDID).
- A configured identifier for the service of the transmitting MEP. This is referred to as the Short MA Name (SMAN). Together, the MDID and the SMAN make up the Maintenance Association Identifier (MAID). The MAID must be configured identically on every MEP in the service.
- These are restrictions on the type of MAID that are supported for sessions with time interval of less than 1 minute. The MAID supports two types of formats on offloaded MEPs:
 - No Domain Name Format
 - MD Name Format = 1-NoDomainName

- Short MA Name Format = 3 - 2 bytes integer value
- Short MA Name Length = 2 - fixed length
- Short MA Name = 2 bytes of integer
- 1731 Maid Format
 - MD Name Format = 1-NoDomainName
 - MA Name Format(MEGID Format) = 32
 - MEGID Length = 13 - fixed length
 - MEGID(ICCCode) = 6 Bytes
 - MEGID(UMC) = 7 Bytes
 - ITU Carrier Code (ICC) - Number of different configurable ICC code - 15 (for each NPU)
 - Unique MEG ID Code (UMC) - 4

Maintenance Association Identifier (MAID) comprises of the Maintenance Domain Identifier (MDID) and Short MA Name (SMAN). MDID only supports **null** value and SMAN only supports ITU Carrier Code (ICC) or a numerical. No other values are supported.

- An example for configuring domain ID null is: **ethernet cfm domain SMB level 3 id null**
- An example for configuring SMAN is: **ethernet cfm domain SMB level 3 id null service 901234AB xconnect group 99999 p2p 99999 id number 1**
- A configured numeric identifier for the MEP (the MEP ID). Each MEP in the service must be configured with a different MEP ID.
- Dynamic Remote MEPs are not supported for MEPs with less than 1 min interval. You must configure MEP CrossCheck for all such MEPs.
- Sequence numbering is not supported for MEPs with less than 1 minute interval.
- In a Remote Defect Indication (RDI), each MEP includes this in the CCMs it is sending, if it has detected a defect relating to the CCMs it is receiving. This notifies all the MEPs in the service that a defect has been detected somewhere in the service.
- The interval at which CCMs are being transmitted.
- CCM Tx/Rx statistics counters are not supported for MEPs with less than 1 minute intervals.
- Sender TLV and Cisco Proprietary TLVs are not supported for MEPs with less than 1 minute intervals.
- The status of the interface where the MEP is operating, for example, whether the interface is up, down, STP blocked, and so on.



Note The status of the interface (up/down) should not be confused with the direction of any MEPs on the interface (Up MEPs/Down MEPs).

These defects can be detected from the received CCMs:

- Interval mismatch: The CCM interval in the received CCM does not match the interval that the MEP is sending CCMs.
- Level mismatch: A MEP has received a CCM carrying a lower maintenance level than the MEP's own level.
- Loop: A CCM is received with the source MAC address equal to the MAC address of the interface where the MEP is operating.
- Configuration error: A CCM is received with the same MEP ID as the MEP ID configured for the receiving MEP.
- Cross-connect: A CCM is received with a MAID that does not match the locally configured MAID. This generally indicates a VLAN misconfiguration within the network, such that CCMs from one service are leaking into a different service.
- Peer interface down: A CCM is received that indicates the interface on the peer is down.
- Remote defect indication: A CCM is received carrying a remote defect indication.



Note This defect does not cause the MEP to include a remote defect indication in the CCMs that it is sending.

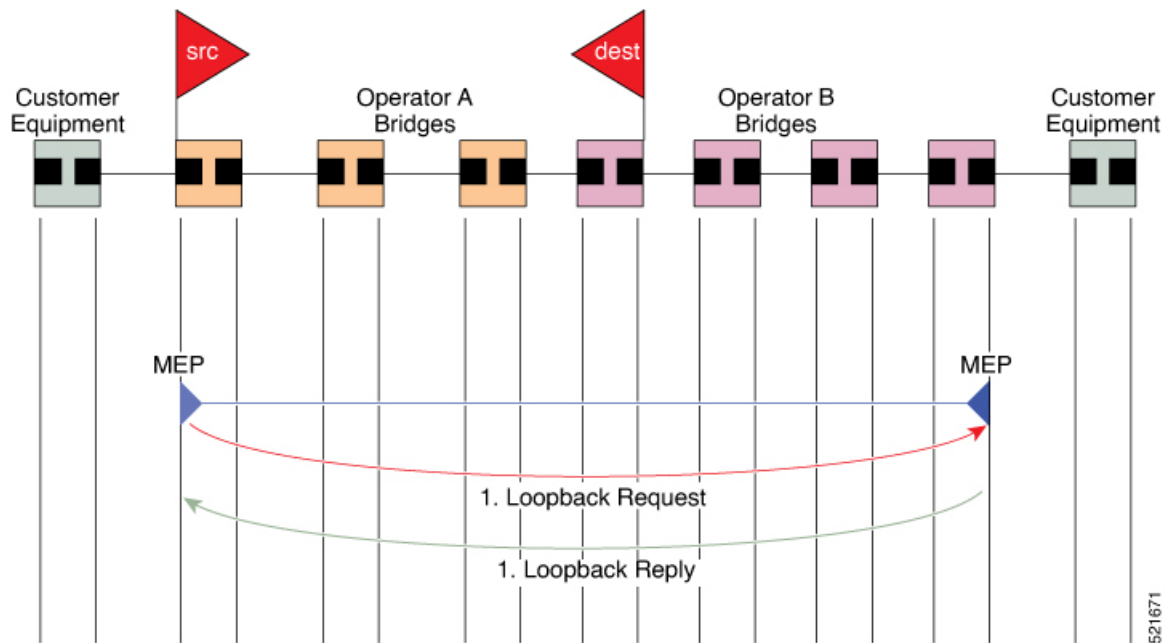
Out-of-sequence CCMs can also be detected by monitoring the sequence number in the received CCMs from each peer MEP. However, this is not considered a CCM defect.

Loopback (IEEE 802.1ag and ITU-T Y.1731)

Loopback Messages (LBM) and Loopback Replies (LBR) are used to verify connectivity between a local MEP and a particular remote MP. At the request of the administrator, a local MEP sends unicast LBMs to the remote MP. On receiving each LBM, the target maintenance point sends an LBR back to the originating MEP. Loopback indicates whether the destination is reachable or not—it does not allow hop-by-hop discovery of the path. It is similar in concept to an ICMP Echo (ping). Since loopback messages are destined for unicast addresses, they are forwarded like normal data traffic, while observing the maintenance levels. At each device that the loopback reaches, if the outgoing interface is known (in the bridge's forwarding database), then the frame is sent out on that interface. If the outgoing interface is not known, then the message is flooded on all interfaces.

This figure shows an example of CFM loopback message flow between a MEP and MEP.

Figure 6: Loopback Messages




Loopback messages can be padded with user-specified data. This allows data corruption to be detected in the network. They also carry a sequence number which allows for out-of-order frames to be detected.

Linktrace (IEEE 802.1ag and ITU-T Y.1731)

Linktrace Messages (LTM) and Linktrace Replies (LTR) are used to track the path (hop-by-hop) to a unicast destination MAC address. At the request of the operator, a local MEP sends an LTM. Each hop where there is a maintenance point sends an LTR back to the originating MEP. This allows the administrator to discover connectivity data about the path. It is similar in concept to IP traceroute, although the mechanism is different. In IP traceroute, successive probes are sent, whereas CFM Linktrace uses a single LTM which is forwarded by each MP in the path. LTMs are multicast, and carry the unicast target MAC address as data within the frame. They are intercepted at each hop where there is a maintenance point, and either retransmitted or dropped to discover the unicast path to the target MAC address.

This figure shows an example of CFM linktrace message flow between MEPs and MEPs.

Figure 7: Linktrace Message Flow



The linktrace mechanism is designed to provide useful information even after a network failure. This allows it to be used to locate failures, for example after a loss of continuity is detected. To achieve this, each MP maintains a CCM Learning Database. This maps the source MAC address for each received CCM to the interface through which the CCM was received. It is similar to a typical bridge MAC learning database, except that it is based only on CCMs and it times out much more slowly—on the order of days rather than minutes.



Note In IEEE 802.1ag, the CCM Learning Database is referred to as the MIP CCM Database. However, it applies to both MIPs and MEPs.

In IEEE 802.1ag, when an MP receives an LTM message, it determines whether to send a reply using the following steps:

1. The target MAC address in the LTM is looked up in the bridge MAC learning table. If the MAC address is known, and therefore the egress interface is known, then an LTR is sent.
2. If the MAC address is not found in the bridge MAC learning table, then it is looked up in the CCM learning database. If it is found, then an LTR is sent.
3. If the MAC address is not found, then no LTR is sent (and the LTM is not forwarded).

If the target MAC has never been seen previously in the network, the linktrace operation will not produce any results.



Note IEEE 802.1ag and ITU-T Y.1731 define slightly different linktrace mechanisms. In particular, the use of the CCM learning database and the algorithm described above for responding to LTM messages are specific to IEEE 802.1ag. IEEE 802.1ag also specifies additional information that can be included in LTRs. Regardless of the differences, the two mechanisms are interoperable.

Configurable Logging

CFM supports logging of various conditions to syslog. Logging can be enabled independently for each service, and when the following conditions occur:

- New peer MEPs are detected, or loss of continuity with a peer MEP occurs.
- Changes to the CCM defect conditions are detected.
- Cross-check “missing” or “unexpected” conditions are detected.
- AIS condition detected (AIS messages received) or cleared (AIS messages no longer received).
- EFD used to shut down an interface, or bring it back up.

How to Configure Ethernet OAM

This section provides these configuration procedures:

Configuring Ethernet OAM

Custom EOAM settings can be configured and shared on multiple interfaces by creating an EOAM profile in Ethernet configuration mode and then attaching the profile to individual interfaces. The profile configuration does not take effect until the profile is attached to an interface. After an EOAM profile is attached to an interface, individual EOAM features can be configured separately on the interface to override the profile settings when desired.

This section describes how to configure an EOAM profile and attach it to an interface in these procedures:

Configuring an Ethernet OAM Profile

Perform these steps to configure an Ethernet OAM profile.

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure terminal	Enters global configuration mode.

	Command or Action	Purpose
Step 2	ethernet oam profile <i>profile-name</i> Example: RP/0/RP0/CPU0:router(config)# ethernet oam profile Profile_1	Creates a new Ethernet Operations, Administration and Maintenance (OAM) profile and enters Ethernet OAM configuration mode.
Step 3	link-monitor Example: RP/0/RP0/CPU0:router(config-eoam)# link-monitor	Enters the Ethernet OAM link monitor configuration mode.
Step 4	symbol-period window { milliseconds <i>window</i> symbols <i>window</i> [thousand million billion] } Example: RP/0/RP0/CPU0:router(config-eoam-lm)# symbol-period window 60000	(Optional) Configures the window size for an Ethernet OAM symbol-period error event. If specified in milliseconds, the range is 1000 to 60000. If not specified as a multiple of 1 second, the actual window used will be rounded up to the nearest second, with thresholds scaled accordingly. If specified in symbols, the range is interface speed dependent (must be between the maximum number of symbols that could be received in 1 second and the maximum number of symbols that could be received in 1 minute). Again the actual window used is rounded up to the nearest second, with thresholds scaled accordingly. The default value is 1000 milliseconds.
Step 5	symbol-period threshold { ppm [low threshold] [high threshold] symbols [low threshold [thousand million billion]] [high threshold [thousand million billion]] } Example: RP/0/RP0/CPU0:router(config-eoam-lm)# symbol-period threshold ppm low 100 high 1000000	(Optional) Configures the thresholds that trigger an Ethernet OAM symbol-period error event, in symbols or ppm (errors per million symbols). When using this command at least one of the high and low thresholds must be specified. If the low threshold is not specified, the default value is used. If the high threshold is not specified, no action is performed in response to an event. The high threshold must not be smaller than the low threshold. If specified in ppm, the range (for both thresholds) is 1 to 1000000. If specified in symbols, the range (for both thresholds) is 1 to the maximum window size in symbols, see Step 4 . The default low threshold is 1 symbol.
Step 6	frame window milliseconds <i>window</i> Example: RP/0/RP0/CPU0:router(config-eoam-lm)# frame window milliseconds 60	(Optional) Configures the frame window size (in milliseconds) of an OAM frame error event. The range is from 1000 to 60000. The default value is 1000.
Step 7	frame threshold [low threshold] [high threshold] Example:	(Optional) Configures the thresholds (in symbols) that triggers an Ethernet OAM frame error event. When using this command at least one of the high and low thresholds

	Command or Action	Purpose
	<pre>RP/0/RP0/CPU0:router(config-eoam-lm)# frame threshold low 10000000 high 60000000</pre>	<p>must be specified. If the low threshold is not specified, the default value is used. If the high threshold is not specified, no action is performed in response to an event. The high threshold must not be smaller than the low threshold.</p> <p>The range is from 1 to 60000000.</p> <p>The default low threshold is 1.</p>
Step 8	<p>frame-period window { milliseconds <i>window</i> frames <i>window</i> [thousand million billion] }</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam-lm)# frame-period window milliseconds 60000</pre>	<p>(Optional) Configures the window size for an Ethernet OAM frame-period error event.</p> <p>The range is from 100 to 60000, if defined in milliseconds. If the window is defined as say, 200ms, and the interface could receive at most say 10000 minimum size frames in 200ms, then the actual window size used will be the time taken to receive 10000 frames, rounded up to the nearest second. The thresholds will be scaled accordingly.</p> <p>If specified in frames, the range is interface speed dependent, but must be between the number of minimum size frames that could be received in 100ms and the number of minimum size frames that could be received in 1 minute. If the window is defined as 20000 frames, the actual window size used will be the time taken to receive 20000 frames, rounded up to the nearest second. The thresholds will be scaled accordingly.</p> <p>The default value is 1000 milliseconds.</p>
Step 9	<p>frame-period threshold { ppm [low threshold] [high threshold] frames [low threshold [thousand million billion]] [high threshold [thousand million billion]] }</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam-lm)# frame-period threshold ppm low 100 high 1000000</pre>	<p>(Optional) Configures the thresholds (either in frames or in ppm - errors per million frames) that trigger an Ethernet OAM frame-period error event. When using this command at least one of the high and low thresholds must be specified. If the low threshold is not specified, the default value is used. If the high threshold is not specified, no action is performed in response to an event. The high threshold must not be smaller than the low threshold.</p> <p>The range for both thresholds is from 1 to 1000000 if specified in ppm. If specified in frames, the range is from 1 to the maximum frame-period window size in frames, see Step 4.</p> <p>The default low threshold is 1 ppm.</p>
Step 10	<p>frame-seconds window milliseconds <i>window</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam-lm)# frame-seconds window milliseconds 900000</pre>	<p>(Optional) Configures the window size (in milliseconds) for the OAM frame-seconds error event.</p> <p>The range is 10000 to 900000.</p> <p>The default value is 6000.</p>
Step 11	<p>frame-seconds threshold [low threshold] [high threshold]</p>	<p>(Optional) Configures the thresholds (in seconds) that trigger a frame-seconds error event. When using this</p>

	Command or Action	Purpose
	<p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam-lm)# frame-seconds threshold low 3 threshold high 900</pre>	<p>command at least one of the high and low thresholds must be specified. If the low threshold is not specified, the default value is used. If the high threshold is not specified, no action is performed in response to an event. The high threshold must not be smaller than the low threshold.</p> <p>The range is 1 to 900</p> <p>The default value is 1.</p>
Step 12	<p>exit</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam-lm)# exit</pre>	Exits back to Ethernet OAM mode.
Step 13	<p>mib-retrieval</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam)# mib-retrieval</pre>	Enables MIB retrieval in an Ethernet OAM profile or on an Ethernet OAM interface.
Step 14	<p>connection timeout <timeout></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam)# connection timeout 30</pre>	<p>Configures the connection timeout period for an Ethernet OAM session. as a multiple of the hello interval.</p> <p>The range is 2 to 30.</p> <p>The default value is 5.</p>
Step 15	<p>hello-interval 1s</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam)# hello-interval 1s</pre>	Configures the time interval between hello packets for an Ethernet OAM session. The default is 1 second (1s).
Step 16	<p>mode {active passive}</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam)# mode passive</pre>	Configures the Ethernet OAM mode. The default is active.
Step 17	<p>require-remote mode {active passive}</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam)# require-remote mode active</pre>	Requires that active mode or passive mode is configured on the remote end before the OAM session becomes active.
Step 18	<p>require-remote mib-retrieval</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-eoam)# require-remote mib-retrieval</pre>	Requires that MIB-retrieval is configured on the remote end before the OAM session becomes active.

	Command or Action	Purpose
Step 19	action capabilities-conflict {disable efd error-disable-interface log} Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action capabilities-conflict efd</pre>	Specifies the action that is taken on an interface when a capabilities-conflict event occurs. The default action is to create a syslog entry.
Step 20	action critical-event {disable error-disable-interface log} Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action critical-event error-disable-interface</pre>	Specifies the action that is taken on an interface when a critical-event notification is received from the remote Ethernet OAM peer. The default action is to create a syslog entry.
Step 21	action discovery-timeout {disable efd error-disable-interface log} Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action discovery-timeout efd</pre>	Specifies the action that is taken on an interface when a connection timeout occurs. The default action is to create a syslog entry.
Step 22	action dying-gasp {disable error-disable-interface log} Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action dying-gasp error-disable-interface</pre>	Specifies the action that is taken on an interface when a dying-gasp notification is received from the remote Ethernet OAM peer. The default action is to create a syslog entry.
Step 23	action high-threshold {disable error-disable-interface log} Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action high-threshold error-disable-interface</pre>	Specifies the action that is taken on an interface when a high threshold is exceeded. The default is to take no action when a high threshold is exceeded.
Step 24	action session-down {disable efd error-disable-interface log} Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action session-down efd</pre>	Specifies the action that is taken on an interface when an Ethernet OAM session goes down.
Step 25	action session-up { disable log } Example: <pre>RP/0/RP0/CPU0:router(config-eoam)# action session-up disable</pre>	Specifies that no action is taken on an interface when an Ethernet OAM session is established. The default action is to create a syslog entry.

	Command or Action	Purpose
Step 26	action uni-directional link-fault {disable efd error-disable-interface log}	Specifies the action that is taken on an interface when a link-fault notification is received from the remote Ethernet OAM peer. The default action is to create a syslog entry. Note In Cisco IOS XR Release 4.x, this command replaces the action link-fault command.
Step 27	action wiring-conflict {disable efd error-disable-interface log} Example: RP/0/RP0/CPU0:router(config-eoam)# action session-down efd	Specifies the action that is taken on an interface when a wiring-conflict event occurs. The default is to put the interface into error-disable state.
Step 28	uni-directional link-fault detection Example: RP/0/RP0/CPU0:router(config-eoam)# uni-directional link-fault detection	Enables detection of a local, unidirectional link fault and sends notification of that fault to an Ethernet OAM peer.
Step 29	commit Example: RP/0/RP0/CPU0:router(config-if)# commit	Saves the configuration changes to the running configuration file and remains within the configuration session.
Step 30	end Example: RP/0/RP0/CPU0:router(config-if)# end	Ends the configuration session and exits to the EXEC mode.

Attaching an Ethernet OAM Profile to an Interface

Perform these steps to attach an Ethernet OAM profile to an interface:

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure terminal	Enters global configuration mode.
Step 2	interface [FastEthernet HundredGigE TenGigE] <i>interface-path-id</i> Example:	Enters interface configuration mode and specifies the Ethernet interface name and notation <i>rack/slot/module/port</i> . Note <ul style="list-style-type: none"> The example indicates an 8-port 10-Gigabit Ethernet interface in modular services card slot 1.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config)# interface TenGigE 0/1/0/0	
Step 3	ethernet oam Example: RP/0/RP0/CPU0:router(config-if)# ethernet oam	Enables Ethernet OAM and enters interface Ethernet OAM configuration mode.
Step 4	profile profile-name Example: RP/0/RP0/CPU0:router(config-if-eoam)# profile Profile_1	Attaches the specified Ethernet OAM profile (<i>profile-name</i>), and all of its configuration, to the interface.
Step 5	commit Example: RP/0/RP0/CPU0:router(config-if)# commit	Saves the configuration changes to the running configuration file and remains within the configuration session.
Step 6	end Example: RP/0/RP0/CPU0:router(config-if)# end	Ends the configuration session and exits to the EXEC mode.

Configuring Ethernet OAM at an Interface and Overriding the Profile Configuration

Using an EOAM profile is an efficient way of configuring multiple interfaces with a common EOAM configuration. However, if you want to use a profile but also change the behavior of certain functions for a particular interface, then you can override the profile configuration. To override certain profile settings that are applied to an interface, you can configure that command in interface Ethernet OAM configuration mode to change the behavior for that interface.

In some cases, only certain keyword options are available in interface Ethernet OAM configuration due to the default settings for the command. For example, without any configuration of the **action** commands, several forms of the command have a default behavior of creating a syslog entry when a profile is created and applied to an interface. Therefore, the **log** keyword is not available in Ethernet OAM configuration for these commands in the profile because it is the default behavior. However, the **log** keyword is available in Interface Ethernet OAM configuration if the default is changed in the profile configuration so you can retain the action of creating a syslog entry for a particular interface.

To see all of the default Ethernet OAM configuration settings, see the [Verifying the CFM Configuration](#).

To configure Ethernet OAM settings at an interface and override the profile configuration, perform these steps:

SUMMARY STEPS

1. **configure**
2. **interface** [HundredGigE | TenGigE] *interface-path-id*
3. **ethernet oam**

4. *interface-Ethernet-OAM-command*
5. **commit**
6. **end**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: <pre>RP/0/RP0/CPU0:router# configure terminal</pre>	Enters global configuration mode.
Step 2	interface [HundredGigE TenGigE] <i>interface-path-id</i> Example: <pre>RP/0/RP0/CPU0:router(config)# interface TenGigE 0/1/0/0</pre>	Enters interface configuration mode and specifies the Ethernet interface name and notation <i>rack/slot/module/port</i> . Note <ul style="list-style-type: none"> • The example indicates an 8-port 10-Gigabit Ethernet interface in modular services card slot 1.
Step 3	ethernet oam Example: <pre>RP/0/RP0/CPU0:router(config-if)# ethernet oam</pre>	Enables Ethernet OAM and enters interface Ethernet OAM configuration mode.
Step 4	<i>interface-Ethernet-OAM-command</i> Example: <pre>RP/0/RP0/CPU0:router(config-if-eoam)# action capabilities-conflict error-disable-interface</pre>	Configures a setting for an Ethernet OAM configuration command and overrides the setting for the profile configuration, where <i>interface-Ethernet-OAM-command</i> is one of the supported commands on the platform in interface Ethernet OAM configuration mode.
Step 5	commit Example: <pre>RP/0/RP0/CPU0:router(config-if)# commit</pre>	Saves the configuration changes to the running configuration file and remains within the configuration session.
Step 6	end Example: <pre>RP/0/RP0/CPU0:router(config-if)# end</pre>	Ends the configuration session and exits to the EXEC mode.

Verifying the Ethernet OAM Configuration

Use the **show ethernet oam configuration** command to display the values for the Ethernet OAM configuration for a particular interface, or for all interfaces. The following example shows the default values for Ethernet OAM settings:

```

RP/0/RP0/CPU0:router# show ethernet oam configuration
Thu Aug  5 22:07:06.870 DST
GigabitEthernet0/4/0/0:
  Hello interval:                        1s
  Mib retrieval enabled:                 N
  Uni-directional link-fault detection enabled: N
  Configured mode:                      Active
  Connection timeout:                   5
  Symbol period window:                 0
  Symbol period low threshold:          1
  Symbol period high threshold:         None
  Frame window:                        1000
  Frame low threshold:                  1
  Frame high threshold:                 None
  Frame period window:                 1000
  Frame period low threshold:           1
  Frame period high threshold:          None
  Frame seconds window:                 60000
  Frame seconds low threshold:          1
  Frame seconds high threshold:         None
  High threshold action:                None
  Link fault action:                    Log
  Dying gasp action:                    Log
  Critical event action:                 Log
  Discovery timeout action:              Log
  Capabilities conflict action:         Log
  Wiring conflict action:                Error-Disable
  Session up action:                     Log
  Session down action:                  Log
  Require remote mode:                  Ignore
  Require remote MIB retrieval:         N

```

Configuring Ethernet CFM

To configure Ethernet CFM, perform the following tasks:



Note CFM is not supported for the following:

- L3 Interfaces and Sub-Interfaces
- Bridge Domain, Release 7.3.1 and earlier
- VPLS, Release 7.3.1 and earlier

Configuring a CFM Maintenance Domain

To configure a CFM maintenance domain, perform the following steps:

SUMMARY STEPS

1. **configure**
2. **ethernet cfm**
3. **domain** *domain-name* **level** *level-value* [**id** [**null**] [**dns** *DNS-name*] [**mac** *H.H.H*] [**string** *string*]]

4. **traceroute cache hold-time** *minutes size entries*
5. **end** or **commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	ethernet cfm Example: RP/0/RP0/CPU0:router(config)# ethernet cfm	Enters Ethernet Connectivity Fault Management (CFM) configuration mode.
Step 3	domain <i>domain-name level level-value [id [null]] [dns DNS-name] [mac H.H.H] [string string]]</i> Example: RP/0/RP0/CPU0:router(config-cfm)# domain Domain_One level 1 id string D1	<p>Creates and names a container for all domain configurations and enters CFM domain configuration mode.</p> <p>The level must be specified.</p> <p>The id is the maintenance domain identifier (MDID) and is used as the first part of the maintenance association identifier (MAID) in CFM frames. If the MDID is not specified, the domain name is used as the MDID by default.</p>
Step 4	traceroute cache hold-time <i>minutes size entries</i> Example: RP/0/RP0/CPU0:router(config-cfm)# traceroute cache hold-time 1 size 3000	(Optional) Sets the maximum limit of traceroute cache entries or the maximum time limit to hold the traceroute cache entries. The default is 100 minutes and 100 entries.
Step 5	end or commit Example: RP/0/RP0/CPU0:router(config-cfm-dmn)# commit	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> • When you use the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> • Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. • Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes.

	Command or Action	Purpose
		<ul style="list-style-type: none"> • Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. • Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring Services for a CFM Maintenance Domain

You can configure up to 50 CFM sessions per line card or 50 CFM sessions per fixed-port router. The system supports 50 CFM sessions on bundles.

Starting Cisco IOS XR Release 7.3.2 and later, 100 CFM sessions are supported for every system.

To configure services for a CFM maintenance domain, perform the following steps:

SUMMARY STEPS

1. **configure**
2. **ethernet cfm**
3. **domain** *domain-name* **level** *level-value* [**id** [null] [**dns** *DNS-name*] [**mac** *H.H.H*] [**string** *string*]]
4. **service** *service-name* {**down-meps** | **xconnect group** *xconnect-group-name* **p2p** *xconnect-name*} [**id** [**icc-based** *icc-string* *umc-string*] | [**number** *number*]
5. **end** or **commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	ethernet cfm Example: RP/0/RP0/CPU0:router(config)# ethernet cfm	Enters Ethernet CFM configuration mode.
Step 3	domain <i>domain-name</i> level <i>level-value</i> [id [null] [dns <i>DNS-name</i>] [mac <i>H.H.H</i>] [string <i>string</i>]] Example: RP/0/RP0/CPU0:router(config-cfm)# domain Domain_One level 1 id string D1	Creates and names a container for all domain configurations at a specified maintenance level, and enters CFM domain configuration mode. The id is the maintenance domain identifier (MDID) and is used as the first part of the maintenance association

	Command or Action	Purpose
		identifier (MAID) in CFM frames. If the MDID is not specified, the domain name is used as the MDID by default.
Step 4	<p>service <i>service-name</i> {down-meps xconnect group <i>xconnect-group-name</i> p2p <i>xconnect-name</i>} [id [icc-based <i>icc-string umc-string</i>] [number <i>number</i>]</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-cfm-dmn) # service xconnect group X1</pre>	<p>Configures and associates a service with the domain and enters CFM domain service configuration mode. You can specify that the service is used only for down MEPs.</p> <p>The id sets the short MA name.</p>
Step 5	<p>end or commit</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-cfm-dmn-svc) # commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> When you use the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]:</pre> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Enabling and Configuring Continuity Check for a CFM Service

To configure Continuity Check for a CFM service, complete the following steps:

SUMMARY STEPS

- configure**
- ethernet cfm**
- domain** *domain-name level level-value* [**id** [**null**] [**dns** *DNS-name*] [**mac** *H.H.H*] [**string** *string*]]
- service** *service-name* {**down-meps** | **xconnect group** *xconnect-group-name* **p2p** *xconnect-name*} [**id** [**icc-based** *icc-string umc-string*] | [**number** *number*]
- continuity-check interval** *time* [**loss-threshold** *threshold*]
- continuity-check archive hold-time** *minutes*
- continuity-check loss auto-traceroute**

8. end or commit

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	ethernet cfm Example: RP/0/RP0/CPU0:router(config)# ethernet cfm	Enters Ethernet Connectivity Fault Management (CFM) configuration mode.
Step 3	domain <i>domain-name</i> level <i>level-value</i> [id [null]] [dns <i>DNS-name</i>] [mac <i>H.H.H</i>] [string <i>string</i>] Example: RP/0/RP0/CPU0:router(config-cfm)# domain Domain_One level 1 id string D1	Creates and names a container for all domain configurations and enters the CFM domain configuration mode. The level must be specified. The id is the maintenance domain identifier (MDID) and is used as the first part of the maintenance association identifier (MAID) in CFM frames. If the MDID is not specified, the domain name is used as the MDID by default.
Step 4	service <i>service-name</i> { down-meps xconnect group <i>xconnect-group-name</i> p2p <i>xconnect-name</i> } [id [icc-based <i>icc-string</i> <i>umc-string</i>] [number <i>number</i>]] Example: RP/0/RP0/CPU0:router(config-cfm-dmn)# service xconnect group X1	Configures and associates a service with the domain and enters CFM domain service configuration mode. You can specify that the service is used only for down MEPs, or associate the service with a xconnect where up MEPs will be created. The id sets the short MA name.
Step 5	continuity-check interval <i>time</i> [loss-threshold <i>threshold</i>] Example: RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# continuity-check interval 100m loss-threshold 10	(Optional) Enables Continuity Check and specifies the time interval at which CCMs are transmitted or to set the threshold limit for when a MEP is declared down.
Step 6	continuity-check archive hold-time <i>minutes</i> Example: RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# continuity-check archive hold-time 100	(Optional) Configures how long information about peer MEPs is stored after they have timed out.
Step 7	continuity-check loss auto-traceroute Example:	(Optional) Configures automatic triggering of a traceroute when a MEP is declared down.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# continuity-check loss auto-traceroute	
Step 8	end or commit Example: RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you use the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]: Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring Cross-Check on a MEP for a CFM Service

To configure cross-check on a MEP for a CFM service and specify the expected set of MEPs, complete the following steps:

SUMMARY STEPS

- configure**
- ethernet cfm**
- domain** *domain-name* **level** *level-value* [**id** [null] [**dns** *DNS-name*] [**mac** *H.H.H*] [**string** *string*]]
- service** *service-name* {**down-meps** | **xconnect group** *xconnect-group-name* **p2p** *xconnect-name*} [**id** [**icc-based** *icc-string umc-string*] | [**string** *text*] | [**number** *number*] | [**vlan-id** *id-number*] | [**vpn-id** *oui-vpnid*]]
- mep crosscheck**
- mep-id** *mep-id-number* [**mac-address** *mac-address*]
- end** or **commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	ethernet cfm Example: RP/0/RP0/CPU0:router# ethernet cfm	Enters the Ethernet Connectivity Fault Management (CFM) configuration mode.
Step 3	domain <i>domain-name</i> level <i>level-value</i> [id [null] [dns <i>DNS-name</i>] [mac <i>H.H.H</i>] [string <i>string</i>]] Example: RP/0/RP0/CPU0:router(config-cfm)# domain Domain_One level 1 id string D1	Creates and names a container for all domain configurations and enters the CFM domain configuration mode. The level must be specified. The id is the maintenance domain identifier (MDID) and is used as the first part of the maintenance association identifier (MAID) in CFM frames. If the MDID is not specified, the domain name is used as the MDID by default.
Step 4	service <i>service-name</i> { down-meps xconnect group <i>xconnect-group-name</i> p2p <i>xconnect-name</i> } [id [icc-based <i>icc-string umc-string</i>] [string <i>text</i>] [number <i>number</i>] [vlan-id <i>id-number</i>] [vpn-id <i>oui-vpnid</i>]]	Configures and associates a service with the domain and enters CFM domain service configuration mode. You can specify that the service is used only for down MEPs, or associate the service with a xconnect where up MEPs will be created. The id sets the short MA name.
Step 5	mep crosscheck Example: RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# mep crosscheck mep-id 10	Enters CFM MEP crosscheck configuration mode.
Step 6	mep-id <i>mep-id-number</i> [mac-address <i>mac-address</i>] Example: RP/0/RP0/CPU0:router(config-cfm-xcheck)# mep-id 10	Enables cross-check on a MEP. Note <ul style="list-style-type: none"> Repeat this command for every MEP that you want included in the expected set of MEPs for cross-check.
Step 7	end or commit Example: RP/0/RP0/CPU0:router(config-cfm-xcheck)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you use the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before</pre>

	Command or Action	Purpose
		<p>exiting (yes/no/cancel) ? [cancel]:</p> <ul style="list-style-type: none"> • Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. • Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. • Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. • Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring Other Options for a CFM Service

To configure other options for a CFM service, complete the following steps:

SUMMARY STEPS

1. **configure**
2. **ethernet cfm**
3. **domain** *domain-name level level-value* [**id** [null] [**dns** *DNS-name*] [**mac** *H.H.H*] [**string** *string*]]
4. **service** *service-name* {**down-meps** | **xconnect group** *xconnect-group-name p2p xconnect-name*} [**id** [**icc-based** *icc-string umc-string*] | [**string** *text*] | [**number** *number*] | [**vlan-id** *id-number*] | [**vpn-id** *oui-vpnid*]]
5. **maximum-meps** *number*
6. **log** {**ais**|**continuity-check errors**|**continuity-check mep changes**|**crosscheck errors**|**efd**}
7. **end** or **commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# <code>configure</code>	Enters global configuration mode.
Step 2	ethernet cfm Example:	Enters the Ethernet Connectivity Fault Management (CFM) configuration mode.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router# ethernet cfm	
Step 3	<p>domain <i>domain-name</i> level <i>level-value</i> [id [null] [dns <i>DNS-name</i>] [mac <i>H.H.H</i>] [string <i>string</i>]]</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-cfm)# domain Domain_One level 1 id string D1</pre>	<p>Creates and names a container for all domain configurations and enters the CFM domain configuration mode.</p> <p>The level must be specified.</p> <p>The id is the maintenance domain identifier (MDID) and is used as the first part of the maintenance association identifier (MAID) in CFM frames. If the MDID is not specified, the domain name is used as the MDID by default.</p>
Step 4	<p>service <i>service-name</i> {down-meps xconnect group <i>xconnect-group-name</i> p2p <i>xconnect-name</i>} [id [icc-based <i>icc-string umc-string</i>] [string <i>text</i>] [number <i>number</i>] [vlan-id <i>id-number</i>] [vpn-id <i>oui-vpnid</i>]]</p>	<p>Configures and associates a service with the domain and enters CFM domain service configuration mode. You can specify that the service is used only for down MEPs, or associate the service with an xconnect where up MEPs will be created.</p> <p>The id sets the short MA name.</p>
Step 5	<p>maximum-meps <i>number</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# maximum-meps 1000</pre>	<p>(Optional) Configures the maximum number (2 to 8190) of MEPs across the network, which limits the number of peer MEPs recorded in the database.</p>
Step 6	<p>log {ais continuity-check errors continuity-check mep changes crosscheck errors efd}</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# log continuity-check errors</pre>	<p>(Optional) Enables logging of certain types of events.</p>
Step 7	<p>end or commit</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> When you use the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes.

	Command or Action	Purpose
		<ul style="list-style-type: none"> Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring CFM MEPs

- For every subinterface configured under a Layer 3 parent interface, you must associate a unique 802.1Q or 802.1ad tag. Else, it leads to unknown network behavior.

SUMMARY STEPS

- configure**
- interface** {**HundredGigE** | **TenGigE**} *interface-path-id*
- interface** {**HundredGigE** | **TenGigE** | **Bundle-Ether**} *interface-path-id* *l2transport*
- ethernet cfm**
- mep domain** *domain-name* **service** *service-name* **mep-id** *id-number*
- cos** *cos*
- end** or **commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# <code>configure</code>	Enters global configuration mode.
Step 2	interface { HundredGigE TenGigE } <i>interface-path-id</i> Example: RP/0/RP0/CPU0:router(config)# <code>interface TenGigE 0/0/0/1</code>	Type of Ethernet interface on which you want to create a MEP. Enter HundredGigE or TenGigE and the physical interface or virtual interface. Note <ul style="list-style-type: none"> Use the show interfaces command to see a list of all interfaces currently configured on the router. L3 interfaces are only supported for bundle member interfaces. Else, you must enable <code>l2transport</code>.
Step 3	interface { HundredGigE TenGigE Bundle-Ether } <i>interface-path-id</i> <i>l2transport</i> Example: RP/0/RP0/CPU0:router(config)# <code>interface TenGigE 0/0/0/1</code>	Type of Ethernet interface on which you want to create a MEP. Enter HundredGigE , TenGigE , or Bundle-Ether and the physical interface or virtual interface followed by the <code>l2transport</code> . <code>l2transport</code> configures the interface as an L2 interface.

	Command or Action	Purpose
		Naming convention is <i>interface-path-id.subinterface</i> . The period in front of the subinterface value is required as part of the notation.
Step 4	ethernet cfm Example: RP/0/RP0/CPU0:router(config-if)# ethernet cfm	Enters interface Ethernet CFM configuration mode.
Step 5	mep domain <i>domain-name</i> service <i>service-name</i> mep-id <i>id-number</i> Example: RP/0/RP0/CPU0:router(config-if-cfm)# mep domain Dm1 service Sv1 mep-id 1	Creates a maintenance end point (MEP) on an interface and enters interface CFM MEP configuration mode.
Step 6	cos <i>cos</i> Example: RP/0/RP0/CPU0:router(config-if-cfm-mep)# cos 7	(Optional) Configures the class of service (CoS) (from 0 to 7) for all CFM packets generated by the MEP on an interface. If not configured, the CoS is inherited from the Ethernet interface. Note For Ethernet interfaces, the CoS is carried as a field in the VLAN tag. Therefore, CoS only applies to interfaces where packets are sent with VLAN tags. If the cos (CFM) command is executed for a MEP on an interface that does not have a VLAN encapsulation configured, it will be ignored.
Step 7	end or commit Example: RP/0/RP0/CPU0:router(config-if-cfm-mep)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you use the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring Y.1731 AIS

This section has the following step procedures:

Configuring AIS in a CFM Domain Service

Use the following procedure to configure Alarm Indication Signal (AIS) transmission for a CFM domain service and configure AIS logging.

SUMMARY STEPS

1. **configure**
2. **ethernet cfm**
3. **domain name level level**
4. **service name xconnect group xconnect-group-name p2p xconnect-name**
5. **ais transmission [interval {1s|1m}][cos cos]**
6. **log ais**
7. **end** or **commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	ethernet cfm Example: RP/0/RP0/CPU0:router(config)# ethernet cfm	Enters Ethernet CFM global configuration mode.
Step 3	domain name level level Example: RP/0/RP0/CPU0:router(config-cfm)# domain D1 level 1	Specifies the domain and domain level.
Step 4	service name xconnect group xconnect-group-name p2p xconnect-name Example: RP/0/RP0/CPU0:router(config-cfm-dmn)# service S1 xconnect group XG1 p2p X2	Specifies the service and cross-connect group and name.
Step 5	ais transmission [interval {1s 1m}][cos cos] Example:	Configures Alarm Indication Signal (AIS) transmission for a Connectivity Fault Management (CFM) domain service.

	Command or Action	Purpose
	RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# ais transmission interval 1m cos 7	
Step 6	log ais Example: RP/0/RP0/CPU0:router(config-cfm-dmn-svc)# log ais	Configures AIS logging for a Connectivity Fault Management (CFM) domain service to indicate when AIS or LCK packets are received.
Step 7	end or commit Example: RP/0/RP0/CPU0:router(config-sla-prof-stat-cfg)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: Uncommitted changes found, commit them before exiting (yes/no/cancel)? [cancel]: Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuring AIS on a CFM Interface

To configure AIS on a CFM interface, perform the following steps:

SUMMARY STEPS

1. **configure**
2. **interface gigabitethernet** *interface-path-id*
3. **ethernet cfm**
4. **ais transmission up interval 1m cos** *cos*
5. **end or commit**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure	Enters global configuration mode.
Step 2	interface gigabitethernet <i>interface-path-id</i> Example: RP/0/RP0/CPU0:router# interface TenGigE 0/0/0/2	Enters interface configuration mode.
Step 3	ethernet cfm Example: RP/0/RP0/CPU0:router(config)# ethernet cfm	Enters Ethernet CFM interface configuration mode.
Step 4	ais transmission up interval 1m cos <i>cos</i> Example: RP/0/RP0/CPU0:router(config-if-cfm)# ais transmission up interval 1m cos 7	Configures Alarm Indication Signal (AIS) transmission on a Connectivity Fault Management (CFM) interface.
Step 5	end or commit Example: RP/0/RP0/CPU0:router(config-sla-prof-stat-cfg)# commit	Saves configuration changes. <ul style="list-style-type: none"> When you issue the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.

Verifying the CFM Configuration

To verify the CFM configuration, use one or more of the following commands:

show ethernet cfm configuration-errors [domain <i>domain-name</i>] [interface <i>interface-path-id</i>]	Displays information about errors that are preventing configured CFM operations from becoming active, as well as any warnings that have occurred.
show ethernet cfm local maintenance-points domain <i>name</i> [service <i>name</i>] interface <i>type interface-path-id</i> [mep mip]	Displays a list of local maintenance points.



Note After you configure CFM, the error message, *cfmd[317]: %L2-CFM-5-CCM_ERROR_CCMS_MISSED : Some received CCMs have not been counted by the CCM error counters*, may display. This error message does not have any functional impact and does not require any action from you.

CFM Over Bundles

CFM over bundle supports the following:

- CFM Maintenance Points — UP MEP, Down MEP, which only includes L2 bundle main and sub-interfaces.
- CCM interval of 100 ms, 1s, 10s, 1min, and 10mins.
- RP OIR/VM reload without impacting learnt CFM peer MEPs.
- Process restart without impacting CFM sessions.
- Static MEPs.

Restrictions for Configuration of CFM on Bundles

Following are the restrictions for configuring CFM over bundle member interfaces:

- Only Layer 2 bundle Ethernet interfaces and sub-interfaces are supported, which are part of a L2VPN cross-connect.
- No support for 3.3ms and 10ms CCM interval.
- Supports 5000 pps rates of CCM traffic for bundle interfaces.
- Ethernet Connectivity Fault Management (CFM) is not supported with Maintenance association End Points (MEPs) that are configured on default and untagged encapsulated sub-interfaces that are part of a single physical interface.
- Multiple MEPs of different directions are not supported on the same interface or Xconnect.
- CFM does not support fast failover, which may result in session flaps on bundle interfaces. Use offload for virtual interfaces to avoid flaps on faster CCM intervals.

Ethernet Frame Delay Measurement for L2VPN Services

Table 18: Feature History Table

Feature Name	Release Information	Feature Description
Ethernet Frame Delay Measurement for L2VPN Services	Release 7.5.3	<p>You can now monitor L2VPN networks and avoid impact to your customers' operations by accurately measuring frame round-trip delays and jitters between two maintenance endpoints (MEPs).</p> <p>This feature lets you detect end-to-end connectivity, loopback, and link trace on MEPs. It reports service performance to your end customers, helping improve technical and operational tasks such as troubleshooting and billing.</p> <p>This feature introduces the cfm-delay-measurement probe command.</p>

Ethernet frame delay measurement complies with the ITU-T Y.1731 standard, which provides comprehensive fault management and performance monitoring recommendations. Delay Measurement Message (DMM) and Delay Measurement Reply (DMR) are used to periodically measure one-way or two-way frame delay and frame delay variation between a pair of point-to-point MEPs. Measurements are made between two MEPs belonging to the same domain and Maintenance Association (MA).

You can measure frame delay in the Layer 2 networks to detect end-to-end connectivity, loopback, and link trace on Maintenance End Points (MEPs) and also report service performance that helps to improve technical and operational tasks such as troubleshooting, billing, and so on. Frame delay is the duration between the time the source node transmits the first bit of a frame and the time the same source node receives the last bit of the frame.

The frame delay measurement uses the following two protocol data units (PDUs):

- Delay Measurement Message (DMM)—DMM is used to measure frame delay and frame delay variation between a pair of point-to-point Maintenance End Points (MEPs).
- Delay Measurement Response (DMR)—DMR is the delay measurement response sent by the destination MEP. When an MEP receives a DMM frame, the responder MEP responds with a DMR frame. The DMR frame carries a reply information and a copy of the timestamp contained in the DMM frame.


We support one-way and two-way frame delay measurement.

Frame Delay Measurement	Description
One-way frame delay measurement (1DM)	<ul style="list-style-type: none"> Measures the frame delay on a unidirectional link between the MEPs. 1DM requires that clocks at both the transmitting MEP and the receiving MEPs are synchronized. Measuring frame-delay variation does not require clock synchronization and the variation can be measured using 1DM and DMR frame combination.
Two-way frame delay measurement	<ul style="list-style-type: none"> Measures the frame delay on a bidirectional link between the MEPs. Two-way delay measurement does not require the clocks at both the transmitting MEP and the receiving MEPs to be synchronized. The two-way frame delay is measured using only DMM and DMR frames.

For more information about CFM, see [Configuring Ethernet OAM, on page 69](#).

Topology

Let's see how a round-trip frame delay is measured with the following sample topology.



- The sender MEP (Router-1) transmits a frame containing delay measurement request information and the timestamp at the which router sends the DMM.
- When packets pass through each interface, timestamps are written into DMMs and DMRs at both local and peer MEPs.
- When the DMM leaves the local interface, the TX timestamp is added to the packet.
- When the receiver MEP (Router-2) receives the frame, records the timestamp at which the receiver MEP receives the frame with the delay measurement request information and the remote MEP (Router-2) responds with an DMR adding the remote TX timestamp to the packet as it leaves the remote interface.

To measure a round-trip delay for a traffic exchange between Router-1 and Router-2, four timestamps get populated as the packet moves through the network.

- Router-1 adds the TxTimestampF when DMM packet is transmitted.
- Router-2 adds RxTimestampF when DMM packet is received by it.

- Router-2 adds TxTimestampB when DMR packet is transmitted.
- Router-1 adds RxTimestampB when DMR is received by it

The round-trip delay is calculated using the following formula:

$$\begin{aligned} \text{Delay} &= (\text{RxTimestampB} - \text{TxTimestampF}) - (\text{TxTimestampB} - \text{RxTimestampF}) \\ &= \text{RxTimestampB} - \text{TxTimestampF} - \text{TxTimestampB} + \text{RxTimestampF} \\ &= (\text{RxTimestampF} - \text{TxTimestampF}) - (\text{TxTimestampB} - \text{RxTimestampB}) \end{aligned}$$

Configure Ethernet Frame Delay Measurement for L2VPN Services

Perform the following tasks to configure Ethernet Frame Delay Measurement for L2VPN Services:

1. Configure L2VPN service.
2. Enable CFM service continuity check.
3. Enable CFM on the interface.
4. Configure Ethernet frame delay measurement.

```

/* Configure L2VPN service */
Router# configure
Router(config)# l2vpn
Router(config-l2vpn)# xconnect group evpn_vpws_203
Router(config-l2vpn-xc)# p2p evpn_vpws_phy-100
Router(config-l2vpn-xc-p2p)# interface GigabitEthernet0/0/0/2.100
Router(config-l2vpn-xc-p2p)# neighbor evpn evi 30001 target 30001 source 50001
Router(config-l2vpn-xc-p2p)# commit

/* Enable CFM service continuity check */
Router# ethernet cfm
Router(config-cfm)# domain xcup1 level 7 id null
Router(config-cfm-dmn)# service xcup1 xconnect group evpn_vpws_Bund
Router(config-cfm-dmn-svc)# mip auto-create all ccm-learning
Router(config-cfm-dmn-svc)# continuity-check interval 1s
Router(config-cfm-dmn-svc)# mep crosscheck
Router(config-cfm-dmn-svc)# mep-id 4001
Router(config-cfm-dmn-svc)# commit

/* Enable CFM on the interface */
Router(config)# interface GigabitEthernet0/0/0/2.100 l2transport
Router(config-subif)# encapsulation dot1q 100
Router(config-subif)# rewrite ingress tag pop 1 symmetric
Router(config-subif)# mtu 9100
Router(config-subif)# ethernet cfm
Router(config-if-cfm)# mep domain bd-domain service bd-service mep-id 4001
Router(config-if-cfm-mep)# sla operation profile test-profile1 target mep-id 1112
Router(config-if-cfm-mep)# commit

/* Configure Ethernet frame delay measurement */
Router(config)# ethernet sla
Router(config-sla)# profile EVC-1 type cfm-delay-measurement
Router(config-sla-prof)# probe
Router(config-sla-prof-pb)# send packet every 1 seconds
Router(config-sla-prof-pb)# schedule
Router(config-sla-prof-schedule)# every 3 minutes for 120 seconds
Router(config-sla-prof-schedule)# statistics
Router(config-sla-prof-stat)# measure round-trip-delay

```

```
Router(config-sla-prof-stat-cfg)# buckets size 1 probes
Router(config-sla-prof-stat-cfg)# buckets archive 5
Router(config-sla-prof-stat-cfg)# commit
```

Running Configuration

This section shows the Ethernet frame delay measurement running configuration.

```
/* Configure L2VPN service */
l2vpn
xconnect group evpn_vpws_203
p2p evpn_vpws_phy-100
interface GigabitEthernet0/0/0/2.100
neighbor evpn evi 30001 target 30001 source 50001
!
/* Enable CFM service continuity check */
ethernet cfm
domain xcup1 level 7 id null
service xcup1 xconnect group evpn_vpws_Bundle_ether203 p2p evpn_vpws-100 id number 4001
mip auto-create all ccm-learning
continuity-check interval 1s
mep crosscheck
mep-id 4001
!
/* Enable CFM on the interface */
interface GigabitEthernet0/0/0/2.100 l2transport
encapsulation dot1q 100
rewrite ingress tag pop 1 symmetric
mtu 9100
ethernet cfm
mep domain bd-domain service bd-service mep-id 4001
sla operation profile test-profile1 target mep-id 1112
!
/* Configure Ethernet SLA */
ethernet sla
profile EVC-1 type cfm-delay-measurement
probe
send packet every 1 seconds
!
schedule
every 3 minutes for 120 seconds
!
statistics
measure round-trip-delay
buckets size 1 probes
buckets archive 5
!
```

Verification

Verify the frame delay measurement. In the following example, you observe that the sent and received DMM and DMR packets are same. So there is no delay in frame transmission.

```
Router# show ethernet cfm local meps interface GigabitEthernet0/0/0/2.100 verbose

Up MEP on GigabitEthernet0/0/0/2.100 MEP-ID 4001
=====
Interface state: Up      MAC address: 0c11.6752.3af8
Peer MEPs: 1 up, 0 with errors, 0 timed out (archived)

CCM generation enabled: Yes, 10s (Remote Defect detected: No)
```

```

AIS generation enabled: No
Sending AIS:           No
Receiving AIS:        No
Sending CSF:          No
Receiving CSF:       No

Packet      Sent      Received
-----
CCM         19       9 (out of seq: 0)
DMM        473         0
DMR         0       473

```

Link Loss Forwarding

Table 19: Feature History Table

Feature Name	Release Information	Feature Description
Link Loss Forwarding	Release 7.9.1	<p>We have now enabled high availability between two bridged interfaces by disabling both interfaces if any one of them fails. Such high availability is enabled because the functionality allows a fault detected on one side of a CFM-protected network to propagate to the other, allowing the device to re-route around the failure.</p> <p>In earlier releases, a failure on one bridged interface did not disable the other interface, and connected devices remained unaware of the link loss.</p> <p>The feature introduces these changes:</p> <ul style="list-style-type: none"> • CLI: New propagate-remote-status command • YANG Data Model: New XPath for Cisco-IOS-XR-um-ethernet-cfm-cfg.yang (see GitHub, YANG Data Models Navigator)

Link Loss Forwarding (LLF) is a mechanism used in networking to propagate the status of a network link to other connected devices. When a link experiences a failure or goes down, LLF ensures that this information is forwarded to other network devices, which can then take appropriate actions to maintain network stability and performance.

You can enable LLF on a network by one of the following methods:

- **Link State Monitor and Propagation by CFM:** LLF uses Connectivity Fault Management (CFM) to transmit notification of a signal loss or fault across the network. When there is a fault on a link to a device on one side of the network, the connection to the port on the other side needs to be shutdown so that the device re-routes the traffic.
- **Remote Link State Propagation:** LLF uses this method for Layer 2 transport events to propagate link failures to remote endpoints. When a link failure occurs, LLF ensures that the failure is communicated to other devices in the network. This enables the other devices to take appropriate action, such as rerouting traffic or triggering failover mechanisms.

Link State Monitor and Propagation by CFM

Link State Monitoring involves tracking the status of network links to ensure they are operational and performing as expected. This can include monitoring for link failures, degradations, or other issues that might affect network performance. When a link state changes, this information needs to be propagated throughout the network so that other devices can adjust their routing tables and network operations accordingly.

When there is a fault on a link to a device on one side of the network, the connection to the port on the other side needs to be shutdown so that the device re-routes the traffic. This requires the interface to be TX-disabled.

Link Loss Forwarding (LLF) uses Connectivity Fault Management (CFM) to transmit notification of a signal loss or fault across the network. If a local attachment circuit (AC) on a bridged interface fails, one of the following signals or packet types are sent to the neighboring device:

- Continuity Check Message (CCM) – The CCMs are heartbeat messages exchanged periodically between all the Maintenance End Points (MEPs) in a service. MEPs are members of a particular service within a domain and are responsible for sourcing and sinking CFM frames. Each MEP sends out multicast CCMs, and receives CCMs from all the other MEPs in the service. This allows each MEP to discover its peer MEPs, and to verify that there is connectivity between them.
- Alarm Indication Signal (AIS) – These are messages sent periodically by MEPs that have detected a fault, to the MEPs in the next highest maintenance domain level.
- Client Signal Fail (CSF) – A mechanism for error detection. When a MEP detects an issue, the MEP sends CSF packets to its peer MEPs.

For more information on MEPs, see [Maintenance Points, on page 78](#).

Connectivity Fault Management Daemon (CFMD) and Ether-MA are processes that run on the control plane of the router. Ether-MA handles owner channel communication and resyncs from CFMD, L2VPN, and other Ether MA processes. This module handles the TX-disable and TX-enable events, based on the notifications from CFMD.

When the system receives a CCM or AIS with fault indication, or a CSF error packet, CFMD communicates with Ether-MA to TX-disable the interface.

When an interface receives a fault notification, the transitions are handled as follows:

- The interface is transitioned to TX-disable state.
- A restore or damping timer with a $3.5 * \text{packet interval}$ duration is started.
- If no other fault packets are received after the restore timer ends, the TX-disable state is cleared and the interface is transitioned to TX-enable state.

Restrictions for Link Loss Forwarding for CFM

- Link loss forwarding is not permitted on subinterfaces.
- Link loss forwarding is permitted only on UP MEPs. The UP MEPs send the frames into the bridge relay function and not through the wire connected to the port where the MEP is configured. For more information on UP MEPs, see [MEP and CFM Processing Overview, on page 79](#).
- A damping or restore timer governs transitions of an interface from TX-disabled state to TX-enabled state. The period of the damping timer is calculated by three times the configured CCM interval. You cannot configure the damping timer.

- The damping timer is not provided for transitions of an interface from TX-enabled state to TX-disabled state.
- Link loss forwarding does not work on bundle interfaces configured with LACP.

Configure Link Loss Forwarding for CFM

To configure LLF on a network:

1. Configure a Connectivity Fault Management (CFM) domain and service.
2. Configure a Maintenance End Point (MEP) under the CFM domain and service.
3. Configure continuity check message (CCM) interval on the MEP. The restore timer for a CCM notification is calculated based on the configured CCM interval.
4. Configure Client Signal Fail (CSF) transmission on the MEP, to enable CSF transmission.
5. Configure CSF logging on the MEP, to enable logging on receiving a CSF packet.



Note The CSF configuration is optional and is not required when both the devices in CFM-protected network are running with IOS-XR. This configuration is required for inter-operation with certain client-end setups that contain devices from other clients.

6. Enable LLF on an interface using the **propagate-remote-status** command. This command triggers the interface to be TX-disabled on fault detection.

Configuration Example

```
/* Configure CFM domain, service, and MEP */
Router# configure
Router(config)# ethernet cfm
Router(config-cfm)# domain dom1 level 1 service ser1 bridge group up-meps bridge-domain
up-mep
/* Configure CCM interval */
Router(config-cfm-dmn-svc)# continuity-check interval 1m
/* (Optional) Configure CSF */
Router(config-cfm-dmn-svc)# csf interval 1m cos 4
Router(config-cfm-dmn-svc)# csf-logging
Router(config-cfm-dmn-svc)# commit
/* Enable LLF on an interface */
Router# configure
Router(config)# interface GigabitEthernet0/2/0/0
Router(config-if)# ethernet cfm
Router(config-if-cfm)# mep domain dom1 service ser1 mep-id 1
Router(config-if-cfm-mep)# propagate-remote-status
Router(config-if-cfm-mep)# commit
```

Running Configuration

```

ethernet cfm
 domain dom1 level 1
  service ser1 bridge group up-meps bridge-domain up-mep
  continuity-check interval 1m
  csf interval 1m cos 4
  csf-logging
  !
!
!
interface GigabitEthernet0/2/0/0
 ethernet cfm
  mep domain dom1 service ser1 mep-id 1
  propagate-remote-status
  !
!
!

```

Verification

The following output shows LLF configuration and fault state for each interface:

```

Router# show ethernet cfm interfaces llf location 0/RP0/CPU0
Defects (from at least one peer MEP):
A - AIS received           I - Wrong interval
R - Remote Defect received V - Wrong Level
L - Loop (our MAC received) T - Timed out (archived)
C - Config (our ID received) M - Missing (cross-check)
X - Cross-connect (wrong MAID) U - Unexpected (cross-check)
P - Peer port down        F - CSF received

```

```

GigabitEthernet0/1/0/0
MEP Defects                               Restore Timer
-----
100 R                                     Not running
101 None                                  10s remaining
102 RPF                                   Not running

```

```

GigabitEthernet0/1/0/1
MEP Defects                               Restore Timer
-----
110 None                                  3s remaining

```

```

GigabitEthernet0/1/0/2
MEP Defects                               Restore Timer
-----
120 P                                     Not running

```

The following output shows that the interface received a single CSF packet at 1 minute interval, so that the interface is TX-disabled with a damping timer of 3.5 minutes.

```

Router# show ethernet cfm local meps detail
Domain dom1 (Level 1), Service ser1
UP MEP on GigabitEthernet0/1/0/0 MEP-ID 1
=====
Interface state: UP    MAC address: 0204.3dbe.c93b
Peer MEPs: 0 up, 0 with errors, 0 timed out (archived)
Cross-check errors: 0 missing, 0 unexpected

CCM generation enabled: No
AIS generation enabled: No
Sending AIS:           No

```

```

Receiving AIS:          No
Sending CSF:           No
Receiving CSF:         Yes (Interval: 1min, started 00:03:29 ago)
TX Disable triggered:  Yes (restore timer not running)

```

The following output shows that the interface received a CCM notification that the peer MEP port is down, so that the interface is TX-disabled.

```

Router# show ethernet cfm local meps detail
Domain dom1 (Level 1), Service ser1
UP MEP on GigabitEthernet0/1/0/0 MEP-ID 1
=====
Interface state: UP      MAC address: 0204.3dbe.c93b
Peer MEPS: 1 up, 1 with errors, 0 timed out (archived)
Cross-check errors: 0 missing, 0 unexpected

CCM generation enabled: Yes, 1min (Remote Defect detected: Yes)
CCM defects detected:   P - peer port down
AIS generation enabled: No
Sending AIS:            No
Receiving AIS:          No
Sending CSF:            No
Receiving CSF:          No
TX Disable triggered:   Yes (restore timer not running)

```

The following output shows that the interface received CCM notification that the peer MEP port is up, and restore timer is started for the TX-disabled interface.

```

Router# show ethernet cfm local meps detail
Domain dom1 (Level 1), Service ser1
UP MEP on GigabitEthernet0/1/0/0 MEP-ID 1
=====
Interface state: UP      MAC address: 0204.3dbe.c93b
Peer MEPS: 1 up, 0 with errors, 0 timed out (archived)
Cross-check errors: 0 missing, 0 unexpected

CCM generation enabled: Yes, 1min (Remote Defect detected: No)
AIS generation enabled: No
Sending AIS:            No
Receiving AIS:          No
Sending CSF:            No
Receiving CSF:          No
TX Disable triggered:   Yes (restore timer running, 1183ms remaining)

```

The following output shows Ether-MA configured bundles and their members:

```

Router# show ethernet infra internal ether-ma bundles
Bundle interface: Bundle-Ether1 (TX disabled)
Bundle members:
  GigabitEthernet0/1/0/1
  GigabitEthernet0/1/0/2

Bundle interface: Bundle-Ether2
Bundle members:
  GigabitEthernet0/2/0/1

```

Remote Link State Propagation

Table 20: Feature History Table

Feature Name	Release Information	Feature Description
Remote Link State Propagation	Release 24.3.1	<p>Introduced in this release on: Fixed Systems (8200, 8700); Centralized Systems (8600); Modular Systems (8800 [LC ASIC: Q100, Q200, P100])</p> <p>Remote Link State Propagation allows the status of a link to be communicated to remote devices, ensuring that all relevant parts of the network are aware of link state changes. Link Loss Forwarding (LLF) uses this feature to propagate link failures to remote endpoints.</p> <p>By enabling remote state propagation and LLF on an interface, you can ensure that the link state changes are communicated to remote devices, allowing for quick failover and rerouting of traffic.</p> <p>This feature introduces the propagate remote-status command.</p>

Remote Link State Propagation allows the status of a link to be communicated to remote devices, ensuring that all relevant parts of the network are aware of link state changes. This is particularly useful in Layer 2 transport networks, where maintaining accurate link status information is crucial for network performance and reliability.

Link Loss Forwarding for Layer 2 Transport

Link Loss Forwarding (LLF) uses Remote Link State Propagation to propagate link failures to remote endpoints. When a link failure occurs, LLF ensures that the failure is communicated to other devices in the network, allowing them to take appropriate action, such as rerouting traffic or triggering failover mechanisms. LLF helps avoid packet loss and triggers network convergence through alternate links. It works by sending signals across the pseudowire (PW) to the neighboring device, bringing the PW and far-end attachment circuit (AC) down if the local AC goes down.

You can configure LLF for the Layer 2 transport events, using the **propagate remote-status** command. You can enable LLF on the following interface types: 1G, 10G, 25G, 40G, 100G, and 400G.

Remote State Propagation with LLF for L2 Transport

When you configure LLF for L2 transport events, the following are the processes that happen:

1. **Link State Detection:** The network device detects a change in the link state, such as a link going down.
2. **Remote State Propagation:** When remote state propagation is enabled, the detected link state change is propagated to remote devices. This ensures that all relevant devices are aware of the link failure and can take appropriate action.
3. **LLF Activation:** LLF uses the propagated link state information to trigger failover mechanisms or reroute traffic. This helps maintain network performance and reliability by quickly responding to link failures.

When you enable remote state propagation and LLF on an interface, the link state changes are communicated to remote devices allowing for quick failover and rerouting of traffic.

Configure Link Loss Forwarding for Layer 2 Transport

The following example shows how to configure LLF for Layer 2 transport events.

Procedure

-
- Step 1** To enable link loss forwarding for Layer transport events, enter the interface mode, configure **l2transport**, and then enable LLF using the **propagate remote-status** command.

```
Router(config)# interface tenGigE 0/0/0/1
Router(config-if)# l2transport
Router(config-if-l2)# propagate remote-status
Router(config-if-l2)# commit
```

- Step 2** View the running configuration.

```
interface TenGigE 0/0/0/1
  l2transport
  propagate remote-status
!
```



CHAPTER 7

IP Event Dampening

The IP Event Dampening feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network operator to configure a router to automatically identify and selectively dampen a local interface that is flapping.

Guidelines and Limitations

See the following guidelines and limitations before configuring IP Event Dampening feature:

- Due to changes in the netstack-IP component, all the IP clients observe the impact of dampening or interface.
 - For each flap of the interface, a certain penalty is added. The penalty decays exponentially whose parameters are configured.
 - When penalty exceeds a certain high level, the interface is dampened. It is unsuppressed when the penalty decays below a low level.
 - When an interface is dampened, the IP address and the static routes are removed from the interface. All the clients of IP get an IP delete notification.
 - When an interface is unsuppressed, the IP address and the relevant routes are added back. All the clients of IP get an IP address add notification for all the IP addresses of the interface.
 - All Layer 3 interfaces that are configured on the Ethernet interface, port changes, and SVI support this feature.
- [IP Event Dampening Overview, on page 121](#)
 - [Interface State Change Events, on page 122](#)
 - [Affected Components, on page 123](#)
 - [How to Configure IP Event Dampening, on page 124](#)

IP Event Dampening Overview

Interface state changes occur when interfaces are administratively brought up or down or if an interface changes state. When an interface changes state or flaps, routing protocols are notified of the status of the routes that are affected by the change in state. Every interface state change requires all affected devices in the network to recalculate best paths, install or remove routes from the routing tables, and then advertise valid routes to peer routers. An unstable interface that flaps excessively can cause other devices in the network to

consume substantial amounts of system processing resources and cause routing protocols to lose synchronisation with the state of the flapping interface.

The IP Event Dampening feature introduces a configurable exponential decay mechanism to suppress the effects of excessive interface flapping events on routing protocols and routing tables in the network. This feature allows the network operator to configure a router to automatically identify and selectively dampen a local interface that is flapping. Dampening an interface removes the interface from the network until the interface stops flapping and becomes stable. Configuring the IP Event Dampening feature improves convergence times and stability throughout the network by isolating failures so that disturbances are not propagated. This, in turn, reduces the utilisation of system processing resources by other devices in the network and improves overall network stability.

Interface State Change Events

This section describes the interface state change events of the IP Event Dampening feature. This feature employs a configurable exponential decay mechanism that is used to suppress the effects of excessive interface flapping or state changes. When the IP Event Dampening feature is enabled, flapping interfaces are dampened from the perspective of the routing protocol by filtering excessive route updates. Flapping interfaces are identified, assigned penalties, suppressed if necessary, and made available to the network when the interface stabilizes.

Suppress Threshold

The suppress threshold is the value of the accumulated penalty that triggers the router to dampen a flapping interface. The flapping interface is identified by the router and assigned a penalty for each up and down state change, but the interface is not automatically dampened. The router tracks the penalties that a flapping interface accumulates. When the accumulated penalty reaches the default or preconfigured suppress threshold, the interface is placed in a dampened state.

Half-Life Period

The half-life period determines how fast the accumulated penalty can decay exponentially. When an interface is placed in a dampened state, the router monitors the interface for additional up and down state changes. If the interface continues to accumulate penalties and the interface remains in the suppress threshold range, the interface will remain dampened. If the interface stabilises and stops flapping, the penalty is reduced by half after each half-life period expires. The accumulated penalty will be reduced until the penalty drops to the reuse threshold. The configurable range of the half-life period timer is from 1 to 45 minutes. The default half-life period timer is 1 minute.

Reuse Threshold

When the accumulated penalty decreases until the penalty drops to the reuse threshold, the route is unsuppressed and made available to other devices in the network. The range of the reuse value is from 1 to 20000 penalties. The default value is 750 penalties.

Maximum Suppress Time

The maximum suppress time represents the maximum time an interface can remain dampened when a penalty is assigned to an interface. The maximum suppress time can be configured from 1 to 255 seconds. The maximum penalty is truncated to maximum 20000 unit. The maximum value of the accumulated penalty is calculated based on the maximum suppress time, reuse threshold, and half-life period.

Affected Components

When an interface is not configured with dampening, or when an interface is configured with dampening but is not suppressed, the routing protocol behavior as a result of interface state transitions is not changed by the IP Event Dampening feature. However, if an interface is suppressed, the routing protocols and routing tables are immune to any further state transitions of the interface until it is unsuppressed.

Route Types

The following interfaces are affected by the configuration of this feature:

- Connected routes:
 - The connected routes of dampened interfaces are not installed into the routing table.
 - When a dampened interface is unsuppressed, the connected routes will be installed into the routing table if the interface is up.
- Static routes:
 - Static routes assigned to a dampened interface are not installed into the routing table.
 - When a dampened interface is unsuppressed, the static route will be installed into the routing table if the interface is up.



Note Only the primary interface can be configured with this feature, and all subinterfaces are subject to the same dampening configuration as the primary interface. IP Event Dampening does not track the flapping of individual subinterfaces on an interface.

Supported Protocols

All the protocols that are used are impacted by the IP Event Dampening feature. The IP Event Dampening feature supports Border Gateway Protocol (BGP), Enhanced Interior Gateway Routing Protocol (EIGRP), Hot Standby Routing Protocol (HSRP), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and VRRP. Ping and SSH to the concerned interface IP address does not work.



Note The IP Event Dampening feature has no effect on any routing protocols if it is not enabled or an interface is not dampened.

How to Configure IP Event Dampening

Enabling IP Event Dampening

The `dampening` command is entered in interface configuration mode to enable the IP Event Dampening feature. If this command is applied to an interface that already has dampening configured, all dampening states are reset and the accumulated penalty will be set to 0. If the interface has been dampened, the accumulated penalty will fall into the reuse threshold range, and the dampened interface will be made available to the network. The flap counts, however, are retained.

Table 21: Procedure

Steps	Command or Action	Purpose
Step 1	configure terminal	Enters global configuration mode.
Step 2	interface <i>type number</i>	Enters interface configuration mode and configures the specified interface.
Step 3	dampening [<i>half-life-period reuse-threshold</i>] [<i>suppress-threshold max-suppress [restart-penalty]</i>]	Enables interface dampening. <ul style="list-style-type: none"> Entering the <code>dampening</code> command without any arguments enables interface dampening with default configuration parameters. When manually configuring the timer for the <code>restart-penalty</code> argument, the values must be manually entered for all arguments.
Step 4	end	Exits interface configuration mode.

Verifying IP Event Dampening

Use the `show dampening interface` or `show interface dampening` commands to verify the configuration of the IP Event Dampening feature.

Table 22: Procedure

Steps	Command or Action	Purpose
Step 1	show dampening interface	Displays dampened interfaces.
Step 2	show interface dampening	Displays dampened interfaces on the local router.



CHAPTER 8

Configuring Link Bundling

Table 23: Feature History Table

Feature Name	Release	Description
1023 Ethernet Bundle Interfaces Support	Release 7.3.15	With the introduction of this enhancement, the maximum system-wide bundle interface scale has increased from 512 to 1023 bundle interfaces. The default value remains at 64-member links for each bundle.
64-bit Bandwidth Support	Release 7.3.15	With this release, the Cisco 8000 Series Router supports 64-bit bandwidth field, as opposed to the previous 32-bit bandwidth field. 64-bit bandwidth enables the system to support interface bandwidths greater than 4.2T.

This module describes the configuration of link bundle interfaces on the Cisco 8000 Series Router.

A link bundle is a group of one or more ports that are aggregated together and treated as a single link.

Each bundle has a single MAC, a single IP address, and a single configuration set (such as ACLs).



Note The router supports both Layer 2 and Layer 3 Link Bundles. If the Link Bundle is a Layer 3 interface, the system requires an IP address. If the Link Bundle is a Layer 2 interface, the system does not require an IP address. A Link Bundle on the router may contain Layer 2 and Layer 3 subinterfaces within it. In which case, the Layer 3 subinterfaces require IP addresses, but the Link Bundle interface does not require an IP address.

The router supports bundling for these types of interfaces:

- Ethernet interfaces

Feature History for Configuring Link Bundling

Release	Modification
Release 7.0.11	Support for this feature added on the router.

Release 7.2.1	Mixed speed bundle members feature added on the router.
---------------	---

- [Limitations and Compatible Characteristics of Ethernet Link Bundles, on page 126](#)
- [Prerequisites for Configuring Link Bundling on a Router, on page 127](#)
- [Information About Configuring Link Bundling, on page 127](#)
- [How to Configure Link Bundling, on page 138](#)
- [Configuration Examples for Link Bundling, on page 147](#)

Limitations and Compatible Characteristics of Ethernet Link Bundles

This list describes the properties and limitations of ethernet link bundles:

- The router supports mixed speed bundles. Mixed speed bundles allow member links of different bandwidth to be configured as active members in a single bundle. The ratio of the bandwidth for bundle members must not exceed 10. Also, the total weight of the bundle must not exceed 64. For example, 100Gbps link and 10Gbps links can be active members in a bundle and load-balancing on member links are based on bandwidth weightage.
- Starting with Cisco IOS XR Release 7.2.1, the router supports mixed speed bundle members in a bundle. Traffic flows are distributed based on the bandwidth of each member link. Mixed speed combinations support a maximum ratio of 10:1. For example, a 10Gbps and a 100Gbps link can be mixed, or a 100Gbps and a 400Gbps link can be mixed, but a 10Gbps and a 400Gbps link cannot be mixed. The following combinations are possible:
 - 400G, 100G
 - 400G, 40G
 - 400G, 100G, 40G
 - 100G, 40G
 - 100G, 10G
 - 100G, 40G, 10G
 - 40G and 10G
- The weight of each bundle member is the ratio of its bandwidth to the lowest bandwidth member. Total weight of the bundle is the sum of weights or relative bandwidth of each bundle member. Since the weight for a bundle member is greater than or equal to 1 and less than or equal to 10, the total member of links in a bundle is less than 64 in mixed bundle case.
- Any type of Ethernet interfaces can be bundled, with or without the use of LACP (Link Aggregation Control Protocol).
- A single router can support a maximum of 512 bundle interfaces and a default of 64 member links per bundle.
-
- With Cisco IOS XR Release 7.3.15, a single router supports up to a maximum of 1023 bundle interfaces and up to a maximum of 64 member links per bundle.

Any newly inserted line card, which leads to exceeding this scale encounters continuous Out of Resource (OOR) failures. To stop the errors, reduce the scale or shut the line card.

- Physical layer and link layer configuration are performed on individual member links of a bundle.
- Configuration of network layer protocols and higher layer applications is performed on the bundle itself.
- IPv4 and IPv6 addressing is supported on ethernet link bundles.
- A bundle can be administratively enabled or disabled.
- Each individual link within a bundle can be administratively enabled or disabled.
- Ethernet link bundles are created in the same way as Ethernet channels, where the user enters the same configuration on both end systems.
- The MAC address that is set on the bundle becomes the MAC address of the links within that bundle.
- Load balancing (the distribution of data between member links) is done by flow instead of by packet. Data is distributed to a link in proportion to the bandwidth of the link in relation to its bundle.
- QoS is supported and is applied proportionally on each bundle member.
- All links within a single bundle must terminate on the same two systems.
- Bundled interfaces are point-to-point.
- A link must be in the up state before it can be in distributing state in a bundle.
- Only physical links can be bundle members.

Prerequisites for Configuring Link Bundling on a Router

Before configuring Link Bundling, ensure that you meet the following tasks and conditions:

- You know the interface IP address (Layer 3 only).
- You know the links that you must include in the bundle that you are configuring.
- If you are configuring an Ethernet link bundle, you must install Ethernet line cards on the router.



Note For more information about physical interfaces, PLIMs, and modular services cards, refer to the *Cisco 8000 Series Router Hardware Installation Guide*.

Information About Configuring Link Bundling

To configure link bundling, you must understand the following concepts:

Link Bundling Overview

The Link Bundling feature allows you to group multiple point-to-point links together into one logical link and provide higher bidirectional bandwidth, redundancy, and load balancing between two routers. The system assigns a virtual interface to the bundled link. You can dynamically add and delete component links from the virtual interface.

The virtual interface is treated as a single interface on which you can configure an IP address and other software features that the link bundle uses. Packets sent to the link bundle are forwarded to one of the links in the bundle.

A link bundle is a group of ports that the system bundles together and the group then acts as a single link. Following are the advantages of link bundles:



Note A single router supports 512 bundles per system.

- Multiple links can span several line cards to form a single interface. Thus, the failure of a single link does not cause a loss of connectivity.
- Bundled interfaces increase bandwidth availability, because traffic is forwarded over all available members of the bundle. Therefore, traffic can flow on the available links, if one of the links within a bundle fails. You can add bandwidth without interrupting the packet flow.

The previous Cisco IOS XR Software Releases stored and processed the interface bandwidth as a 32-bit value, which supported bundles with an aggregate bandwidth (from the sum of its members) up to 4.2 Gbps. With Cisco IOS XR Software Release 7.3.15, the interface bandwidth is stored and processed as a 64-bit value. The 64-bit value supports larger aggregate bandwidth as some bundles now contain sufficient high bandwidth members that the aggregate value can exceed the 4.2 Gbps bandwidth.

All the individual links within a single bundle must be of the same type.

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

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

Link Aggregation Through LACP

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

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

LACP transmits frames containing the local port state and the local view of the partner system's state. The system analyzes these frames to ensure that both the systems are in agreement.

IEEE 802.3ad Standard

The IEEE 802.3ad standard typically defines a method of forming Ethernet link bundles.

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

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

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

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

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

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



Note We recommend that you avoid modifying the MAC address, because changes in the MAC address can affect packet forwarding.

Configuring LACP Fallback

This section describes how to configure the LACP Fallback feature.

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface Bundle-Ether *bundle-id***

Example:

```
RP/0/RP0/CPU0:router(config)# interface Bundle-Ether 3
```

Creates and names a new Ethernet link bundle.

This **interface Bundle-Ether** command enters you into the interface configuration submode, where you can enter interface-specific configuration commands. Use the **exit** command to exit from the interface configuration submode back to the normal global configuration mode.

Step 3 **ipv4 address** *ipv4-address mask*

Example:

```
RP/0/RP0/CPU0:router(config-if)# ipv4 address 192.168.1.27 255.0.0.0
```

Specifies a primary IPv4 address for an interface.

Step 4 **bundle lacp-fallback timeout 4** *number*

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle lacp-fallback timeout 4
```

Enables the LACP Fallback feature.

Step 5 **end** or **commit**

Example:

```
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

Step 6 **show bundle infrastructure database ma bdl-info Bundle-e1010 | inc** *text*

Example:

```
RP/0/RP0/CPU0:router# show bundle infrastructure database ma bdl-info Bundle-e1010 | inc "fallback"
```

(Optional) Shows the MA information of the bundle manager.

Step 7 **show bundle infrastructure database ma bdl-info Bundle-e1015 | inc** *text*

Example:

```
RP/0/RP0/CPU0:router# show bundle infrastructure database ma bdl-info Bundle-e1015 | inc "fallback"
```

(Optional) Shows the MA information of the bundle manager.

LACP Short Period Time Intervals

As packets are exchanged across member links of a bundled interface, some member links may slow down or time-out and fail. LACP packets are exchanged periodically across these links to verify the stability and reliability of the links over which they pass. The configuration of short period time intervals, in which LACP packets are sent, enables faster detection and recovery from link failures.

Short period time intervals are configured as follows:

- In milliseconds

- In increments of 100 milliseconds
- In the range 100 to 1000 milliseconds
- The default is 1000 milliseconds (1 second)
- Up to 64 member links
- Up to 1280 packets per second (pps)

After 6 missed packets, the link is detached from the bundle.

When the short period time interval is *not* configured, LACP packets are transmitted over a member link every 30 seconds by default.

When the short period time interval is configured, LACP packets are transmitted over a member link once every 1000 milliseconds (1 second) by default. Optionally, both the transmit and receive intervals can be configured to less than 1000 milliseconds, independently or together, in increments of 100 milliseconds (100, 200, 300, and so on).

When you configure a custom LACP short period *transmit* interval at one end of a link, you must configure the same time period for the *receive* interval at the other end of the link.



Note You must always configure the *transmit* interval at both ends of the connection before you configure the *receive* interval at either end of the connection. Failure to configure the *transmit* interval at both ends first results in route flapping (a route going up and down continuously). When you remove a custom LACP short period, you must do it in reverse order. You must remove the *receive* intervals first and then the *transmit* intervals.

Load Balancing

Load balancing is a forwarding mechanism that distributes traffic over multiple links that are based on certain parameters. The router support load balancing for all links in a bundle using Layer 2, Layer 3, and Layer 4 routing information. Starting with Cisco IOS XR Software Release 7.2.1, bandwidth based load-balancing is applicable to L3 unicast flows.

This section describes the load balancing support on link bundles.

For more information about other forms of load balancing on the router, see the following:

- Per-flow load balancing on non-bundle interfaces using Layer 3 and 4 routing information.
- Pseudowire (PW) Load Balancing beginning in Cisco IOS XR 4.0.1.

Layer 3 Egress Load Balancing on Link Bundles

Layer 3 load balancing support began on the router in Cisco IOS XR 7.0.11 release.

Layer 3 load balancing for link bundles is enabled globally by default.

The ingress linecard does bundle member selection and forwards the packet to the linecard and network processor (NP) corresponding to the selected bundle member. The same hash value is used for both ingress and egress linecards. Therefore, even though the egress linecard also does bundle member selection, it selects the same bundle member that was selected by the ingress linecard.

Multicast IPv4 and IPv6 Traffic

For outbound multicast IPv4 or IPv6 traffic, a set of egress linecards is predetermined by the system. If a bundle interface or bundle subinterface is an outgoing interface, the system selects the bundle member for each outgoing interface in a route based on the multicast group address. This helps with load distribution of multicast routed traffic to different bundle members, while providing traffic sequencing within a specific route.

The egress linecard does NP selection using the same approach, when bundle members are spread across multiple NPs within the egress linecard.

When the packet arrives on an egress NP, it uses the 5-tuple hash to select a bundle member within an NP for each packet. This provides better resiliency for bundle member state changes within an NP.

Configuring the Default LACP Short Period Time Interval

This section describes how to configure the default short period time interval for sending and receiving LACP packets on a Gigabit Ethernet interface. This procedure also enables the LACP short period.

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface HundredGigEinterface-path**

Example:

```
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/1
```

Creates a Gigabit Ethernet interface and enters interface configuration mode.

Step 3 **bundle id number mode active**

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle id 1 mode active
```

Specifies the bundle interface and puts the member interface in active mode.

Step 4 **lacp period short**

Example:

```
RP/0/RP0/CPU0:router(config-if)# lacp period short
```

Configures a short period time interval for the sending and receiving of LACP packets, using the default time period of 1000 milliseconds or 1 second.

Step 5 **end or commit**

Example:

```
RP/0/RP0/CPU0:router(config-if)# end
```

or

```
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before
exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.

- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.

- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Example

This example shows how to configure the LACP short period time interval to the default time of 1000 milliseconds (1 second):

```
config
interface HundredGigE 0/1/0/1
  bundle id 1 mode active
  lacp period short
  commit
```

Configuring Custom LACP Short Period Time Intervals

This section describes how to configure custom short period time interval for sending and receiving LACP packets on a Gigabit Ethernet interface.

Procedure

Step 1 **configure****Example:**

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface HundredGigE***interface-path*

Example:

```
RP/0/RP0/CPU0:router(config)# interface HundredGigE 0/1/0/1
```

Creates a Gigabit Ethernet interface and enters interface configuration mode.

Step 3 **bundle id** *number* **mode active**

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle id 1 mode active
```

Specifies the bundle interface and puts the member interface in active mode.

Step 4 **lacp period***time-interval*

Example:

```
RP/0/RP0/CPU0:router(config-if)# lacp period 300
```

Configures a custom period time interval for the sending and receiving of LACP packets. The interval can be in the range 100 to 1000 ms, in multiples of 100.

Step 5 **end** or **commit**

Example:

```
RP/0/RP0/CPU0:router(config-if)# end
```

or

```
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before
exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.

- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.

- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Example

This example shows how to configure the LACP period time interval to the custom time of 300 milliseconds:

```
config
interface HundredGigE 0/1/0/1
  bundle id 1 mode active
  lacp period 300
commit
```

QoS and Link Bundling

On the router, when the system applies QoS on the bundle for either the ingress or egress direction, QoS is applied at each member interface. For complete information on configuring QoS on link bundles on the router, refer to the *Cisco 8000 Series Aggregation Services Router Modular Quality of Service Configuration Guide* and the *Cisco 8000 Series Aggregation Services Router Modular Quality of Service Command Reference*.

Link Bundle Configuration Overview

The following steps provide a general overview of the link bundle configuration process. Ensure that you clear all previous network layer configuration before adding it to a bundle:

1. In global configuration mode, create a link bundle. To create an Ethernet link bundle, enter the **interface Bundle-Ether** command.
2. Assign an IP address and subnet mask to the virtual interface using the **ipv4 address** command.
3. Add interfaces to the bundle that you created in Step 1 with the **bundle id** command in the interface configuration submode.

You can add up to 64 links to a single bundle.



Note The system configures a link as a member of a bundle from the interface configuration submode for that link.

Nonstop Forwarding During Card Failover

Cisco IOS XR software supports nonstop forwarding during a failover between active and standby paired RP cards. Nonstop forwarding ensures that there is no change in the state of the link bundles when a failover occurs.

For example, if an active RP fails, the standby RP becomes operational. The system replicates the configuration, node state, and checkpoint data of the failed RP to the standby RP. The bundled interfaces are present when the standby RP becomes the active RP.



Note Failover is always onto the standby RP.

You do not need to configure anything to guarantee that the system maintains the standby interface configurations.

Link Failover

When one member link in a bundle fails, the system redirects the traffic to the remaining operational member links and traffic flow remains uninterrupted.

Link Switchover

By default, a maximum of 64 links in a bundle can actively carry traffic. If one member link in a bundle fails, traffic is redirected to the remaining operational member links.

You can optionally implement 1:1 link protection for a bundle by setting the **bundle maximum-active links** command to 1. By doing so, you designate one active link and one or more dedicated standby links. If the active link fails, a switchover occurs and a standby link immediately becomes active, thereby ensuring uninterrupted traffic.

If the active and standby links are running LACP, you can choose between an IEEE standard-based switchover (the default) or a faster proprietary optimized switchover. If the active and standby links are not running LACP, the proprietary optimized switchover option is used.

Regardless of the type of switchover you are using, you can disable the wait-while timer, which expedites the state negotiations of the standby link and causes a faster switchover from a failed active link to the standby link.

To do so, you can use the **lacp switchover suppress-flaps** command.

LACP Fallback

The LACP Fallback feature allows an active LACP interface to establish a Link Aggregation Group (LAG) port-channel before the port-channel receives the Link Aggregation and Control Protocol (LACP) protocol data units (PDU) from its peer.

With the LACP Fallback feature configured, the router allows the server to bring up the LAG, before receiving any LACP PDUs from the server, and keeps one port active. This allows the server to establish a connection to PXE server over one Ethernet port, download its boot image and then continue the booting process. When the server boot process is complete, the server fully forms an LACP port-channel.

Designate a Member Link as Unviable

Table 24: Feature History Table

Feature Name	Release	Description
Designate a Member Link as Unviable	Release 7.10.1	<p>Earlier, when a member link is added to an interface link bundle, the peer starts using the link as soon as the LACP communication is up. Sometimes, the hardware programming for the data-path does not get complete in this time resulting in packet loss without any notification to the source.</p> <p>You can now mark a member link as unviable to introduce a delay during which the link is treated as standby. By delaying the usage of the member link for data transmission, you can ensure that the link configuration is fully established, which enables successful data transmission.</p> <p>This feature introduces these changes:</p> <ul style="list-style-type: none"> • CLI: forwarding-unviable • YANG Data Model: New XPath for <code>Cisco-IOS-XR-bundlemgr-oper</code> (see Github, YANG Data Models Navigator).

Link bundling aggregates multiple physical links to a single logical link. When a member link is added to a link bundle, link aggregation control protocol (LACP) communication gets established with the peer to negotiate and control the link aggregation. LACP doesn't have any provision to incorporate a delay before letting data transmission over the link. Therefore, the peer starts using the link when the LACP communication is up. Occasionally, even though the link status is up, and LACP communication is up, the hardware programming for data-path packet forwarding doesn't get complete. In such scenarios, the transmitted data gets lost without any notification or error message to the source or destination of the traffic.

You can now delay the use of such member links, which aren't fully ready to handle data transmission, using the **forwarding-unviable** command. This command configures the link as forwarding-unviable and the member link is considered "standby" for bundle management. As standby member links of a bundle aren't used for data transmission, the usage of forwarding-unviable member links is delayed. When the member link is fully up, that is, the packet forwarding data-path is also completely programmed, you can disable forwarding-unviability of the link using **no forwarding-unviable** command. This removes the forwarding-unviable configuration of the link. Then, the link is treated as an "active" member of the bundle and is used in data transmission and load balancing.



Note It is recommended to wait for a few minutes before running **no forwarding-unviable** command to ensure that the packet forwarding data-path is completely programmed.

Guidelines and Restrictions for Designating Member Links as Unviable

- Forwarding-unviable is disabled on all Ethernet interfaces by default. Therefore, by default, all member links in a bundle are considered "active".
- A link bundle is considered up, only if at least one member link is active. Only the active member links in the link bundle are used for data transmission, load balancing, and redundancy.
- If a link bundle has only one member link, which is forwarding-unviable, the bundle state is considered "down".
- If all the member links in a bundle are forwarding-unviable, the bundle state is considered "down".
- Other existing threshold parameters such as minimum-active links, maximum-active links, and maximum-active-bandwidth, which are considered to determine the bundle state, continue to function along with forwarding-unviable functionality. For more details on these parameters, see [How to Configure Link Bundling, on page 138](#).
- There is no effect of forwarding-unviable configuration on individual Ethernet interfaces that are not part of a link bundle. That is, irrespective of the configuration, such non-member interfaces continue to attempt data transmission and reception.

How to Configure Link Bundling

This section contains the following procedures:

Configuring Ethernet Link Bundles

This section describes how to configure an Ethernet link bundle.



Note In order for an Ethernet bundle to be active, you must perform the same configuration on both connection endpoints of the bundle.



Tip You can programmatically perform the configuration using `openconfig-lacp.yang`, `openconfig-if-aggregate.yang` OpenConfig data models, `Cisco-IOS-XR-bundlemgr-oper.yang` Cisco IOS XR native data model or `Cisco-IOS-XR-um-lacp-cfg.yang` Unified data model. To get started with using data models, see the *Programmability Configuration Guide for Cisco 8000 Series Routers*.

Procedure

Step 1 configure

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface Bundle-Ether** *bundle-id*

Example:

```
RP/0/RP0/CPU0:router#(config)# interface Bundle-Ether 3
```

Creates a new Ethernet link bundle with the specified bundle-id. The range is 1 to 65535.

This **interface Bundle-Ether** command enters you into the interface configuration submode, where you can enter interface specific configuration commands are entered. Use the **exit** command to exit from the interface configuration submode back to the normal global configuration mode.

Step 3 **ipv4 address** *ipv4-address mask*

Example:

```
RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.1.2.3 255.0.0.0
```

Assigns an IP address and subnet mask to the virtual interface using the **ipv4 address** configuration subcommand.

Note • On the router, only a Layer 3 bundle interface requires an IP address.

Step 4 **bundle minimum-active bandwidth** *kbps*

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle minimum-active bandwidth 580000
```

(Optional) Sets the minimum amount of bandwidth required before a user can bring up a bundle.

Step 5 **bundle minimum-active links** *links*

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle minimum-active links 2
```

(Optional) Sets the number of active links required before you can bring up a specific bundle.

Step 6 **bundle maximum-active links** *links* [**hot-standby**]

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle maximum-active links 1 hot-standby
```

(Optional) Implements 1:1 link protection for the bundle, which causes the highest-priority link in the bundle to become active and the second-highest-priority link to become the standby. Also, specifies that a switchover between active and standby LACP-enabled links is implemented per a proprietary optimization.

Note • The priority of the active and standby links is based on the value of the **bundle port-priority** command.

Step 7 **lACP fast-switchover**

Example:

```
RP/0/RP0/CPU0:router(config-if)# lACP fast-switchover
```

(Optional) If you enabled 1:1 link protection (you set the value of the **bundle maximum-active links** command to 1) on a bundle with member links running LACP, you can optionally disable the wait-while timer in the LACP state machine. Disabling this timer causes a bundle member link in standby mode to expedite its normal state negotiations, thereby enabling a faster switchover from a failed active link to the standby link.

Step 8 **exit**

Example:

```
RP/0/RP0/CPU0:router(config-if)# exit
```

Exits interface configuration submode for the Ethernet link bundle.

Step 9 **interface {GigabitEthernet | TenGigE} interface-path-id**

Example:

```
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 1/0/0/0
```

Enters interface configuration mode for the specified interface.

Enter the **GigabitEthernet** or **TenGigE** keyword to specify the interface type. Replace the *interface-path-id* argument with the node-id in the *rack/slot/module* format.

Step 10 **bundle id bundle-id [mode {active | on | passive}]**

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle-id 3
```

Adds the link to the specified bundle.

To enable active or passive LACP on the bundle, include the optional **mode active** or **mode passive** keywords in the command string.

To add the link to the bundle without LACP support, include the optional **mode on** keywords with the command string.

Note • If you do not specify the **mode** keyword, the default mode is **on** (LACP is not run over the port).

Step 11 **bundle port-priority priority**

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle port-priority 1
```

(Optional) If you set the **bundle maximum-active links** command to 1, you must also set the priority of the active link to the highest priority (lowest value) and the standby link to the second-highest priority (next lowest value). For example, you can set the priority of the active link to 1 and the standby link to 2.

Step 12 **no shutdown**

Example:

```
RP/0/RP0/CPU0:router(config-if)# no shutdown
```

(Optional) If a link is in the down state, bring it up. The **no shutdown** command returns the link to an up or down state depending on the configuration and state of the link.

Step 13 **exit****Example:**

```
RP/0/RP0/CPU0:router(config-if)# exit
```

Exits interface configuration submode for the Ethernet interface.

Step 14 **bundle id** *bundle-id* [**mode** {**active** | **passive** | **on**}] **no shutdown** **exit****Example:**

```
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 1/0/2/1
```

```
RP/0/RP0/CPU0:router(config-if)# bundle id 3
```

```
RP/0/RP0/CPU0:router(config-if)# bundle port-priority 2
```

```
RP/0/RP0/CPU0:router(config-if)# no shutdown
```

```
RP/0/RP0/CPU0:router(config-if)# exit
```

```
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 1/0/2/3
```

```
RP/0/RP0/CPU0:router(config-if)# bundle id 3
```

```
RP/0/RP0/CPU0:router(config-if)# no shutdown
```

```
RP/0/RP0/CPU0:router(config-if)# exit
```

(Optional) Repeat Step 8 through Step 11 to add more links to the bundle.

Step 15 **end** or **commit****Example:**

```
RP/0/RP0/CPU0:router(config-if)# end
```

or

```
RP/0/RP0/CPU0:router(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?  
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Step 16 **exit****Example:**

```
RP/0/RP0/CPU0:router(config-if)# exit
```

Exits interface configuration mode.

Step 17 **exit****Example:**

```
RP/0/RP0/CPU0:router(config)# exit
```

Exits global configuration mode.

Step 18 Perform Step 1 through Step 15 on the remote end of the connection.

Brings up the other end of the link bundle.

Step 19 **show bundle Bundle-Ether *bundle-id*****Example:**

```
RP/0/RP0/CPU0:router# show bundle Bundle-Ether 3
```

(Optional) Shows information about the specified Ethernet link bundle.

Step 20 **show lacp bundle Bundle-Ether *bundle-id*****Example:**

```
RP/0/RP0/CPU0:router# show lacp bundle  
Bundle-Ether 3
```

(Optional) Shows detailed information about LACP ports and their peers.

Configuring VLAN Bundles

This section describes how to configure a VLAN bundle. The creation of a VLAN bundle involves three main tasks:

1. Create an Ethernet bundle.
2. Create VLAN subinterfaces and assign them to the Ethernet bundle.
3. Assign Ethernet links to the Ethernet bundle.

These tasks are described in detail in the procedure that follows.



Note In order for a VLAN bundle to be active, you must perform the same configuration on both ends of the bundle connection.

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0:router# configure
```

Enters global configuration mode.

Step 2 **interface Bundle-Ether** *bundle-id*

Example:

```
RP/0/RP0/CPU0:router#(config)# interface Bundle-Ether 3
```

Creates and names a new Ethernet link bundle.

This **interface Bundle-Ether** command enters you into the interface configuration submenu, where you can enter interface-specific configuration commands. Use the **exit** command to exit from the interface configuration submenu back to the normal global configuration mode.

Step 3 **ipv4 address** *ipv4-address mask*

Example:

```
RP/0/RP0/CPU0:router(config-if)# ipv4 address 10.1.2.3 255.0.0.0
```

Assigns an IP address and subnet mask to the virtual interface using the **ipv4 address** configuration subcommand.

Step 4 **bundle minimum-active links** *links*

Example:

```
RP/0/RP0/CPU0:router(config-if)# bundle minimum-active links 2
```

(Optional) Sets the number of active links required before you can bring up a specific bundle.

Step 5 **exit**

Example:

```
RP/0/RP0/CPU0:router(config-if)# exit
```

Exits the interface configuration submenu.

Step 6 **interface Bundle-Ether** *bundle-id.vlan-id*

Example:

```
RP/0/RP0/CPU0:router#(config)# interface Bundle-Ether 3.1
```

Creates a new VLAN, and assigns the VLAN to the Ethernet bundle you created in Step 2.

Replace the *bundle-id* argument with the *bundle-id* you created in Step 2.

Replace the *vlan-id* with a subinterface identifier. Range is from 1 to 4094 inclusive (0 and 4095 are reserved).

Note When you include the *.vlan-id* argument with the **interface Bundle-Ether** *bundle-id* command, you enter subinterface configuration mode.

Step 7 **encapsulation dot1q**

Example:

```
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 100
```

Sets the Layer 2 encapsulation of an interface.

Step 8 **ipv4 address** *ipv4-address mask*

Example:

```
RP/0/RP0/CPU0:router#(config-subif)# ipv4 address 10.1.2.3/24
```

Assigns an IP address and subnet mask to the subinterface.

Step 9 **no shutdown**

Example:

```
RP/0/RP0/CPU0:router#(config-subif)# no shutdown
```

(Optional) If a link is in the down state, bring it up. The **no shutdown** command returns the link to an up or down state depending on the configuration and state of the link.

Step 10 **exit**

Example:

```
RP/0/RP0`/CPU0:router(config-subif)# exit
```

Exits subinterface configuration mode for the VLAN subinterface.

Step 11 Repeat Step 9 through Step 12 to add more VLANs to the bundle you created in Step 2.

(Optional) Adds more subinterfaces to the bundle.

Step 12 **end** or **commit**

Example:

```
RP/0/RP0/CPU0:router(config-subif)# end
```

or

```
RP/0/RP0/CPU0:router(config-subif)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before
exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Step 13 **exit****Example:**

```
RP/0/RP0/CPU0:router(config-subif)# end
```

Exits interface configuration mode.

Step 14 **exit****Example:**

```
RP/0/RP0/CPU0:router(config)# exit
```

Exits global configuration mode.

Step 15 **configure****Example:**

```
RP/0/RP0/CPU0:router # configure
```

Enters global configuration mode.

Step 16 **interface {GigabitEthernet | TenGigE} interface-path-id****Example:**

```
RP/0/RP0/CPU0:router(config)# interface GigabitEthernet 1/0/0/0
```

Enters interface configuration mode for the Ethernet interface you want to add to the Bundle.

Enter the **GigabitEthernet** or **TenGigE** keyword to specify the interface type. Replace the *interface-path-id* argument with the node-id in the rack/slot/module format.

Note A VLAN bundle is not active until you add an Ethernet interface on both ends of the link bundle.

Step 17 **lACP fast-switchover****Example:**

```
RP/0/RP0/CPU0:router(config-if)# lacp fast-switchover
```

(Optional) If you enabled 1:1 link protection (you set the value of the **bundle maximum-active links** command to 1) on a bundle with member links running LACP, you can optionally disable the wait-while timer in the LACP state machine. Disabling this timer causes a bundle member link in standby mode to expedite its normal state negotiations, thereby enabling a faster switchover from a failed active link to the standby link.

VLANs on an Ethernet Link Bundle

You can configure 802.1Q VLAN subinterfaces on 802.3ad Ethernet link bundles.



Note The memory requirement for bundle VLANs is slightly higher than standard physical interfaces.

To create a VLAN subinterface on a bundle, include the VLAN subinterface instance with the **interface Bundle-Ether** command, as follows:

```
interface Bundle-Ether interface-bundle-id.subinterface
```

After you create a VLAN on an Ethernet link bundle, the system supports all VLAN subinterface configuration on that link bundle.

VLAN subinterfaces can support Ethernet Flow Points (EFPs) and Layer 3 services.

You can configure Layer 3 VLAN subinterfaces as follows:

```
interface bundle-ether instance.subinterface, encapsulation dot1q xxxxx
```

Configuring a Member Link as Unviable

Perform the following task to designate a member link as unviable.

Example Configuration

```
RP/0/RP0/CPU0:ios (config) #interface HundredGigE 0/0/0/34
RP/0/RP0/CPU0:ios (config-if) #forwarding-unviable
RP/0/RP0/CPU0:ios (config-if) #commit
RP/0/RP0/CPU0:ios (config-if) #end
```

Running Configuration

```
RP/0/RP0/CPU0:ios#show running-config interface HundredGigE 0/0/0/34
Thu Apr 20 11:11:55.744 UTC
interface HundredGigE0/0/0/34
  forwarding-unviable
!
```

Verification

Use **show bundle** command to view the forwarding-viable status of LAG members. Here, the interface HundredGigE 0/0/0/34 is added to ethernet bundle 3.

```
RP/0/RP0/CPU0:ios#show bundle
Thu Apr 20 11:29:42.500 UTC

Bundle-Ether3
  Status: Down
  Local links <active/standby/configured>: 0 / 0 / 1
  Local bandwidth <effective/available>: 0 (0) kbps
  MAC address (source): 78c6.9991.3504 (Chassis pool)
  Inter-chassis link: No
  Minimum active links / bandwidth: 1 / 1 kbps
  Maximum active links: 64
  Wait while timer: 2000 ms
```

```

Load balancing:
  Link order signaling:      Not configured
  Hash type:                 Default
  Locality threshold:       None
LACP:
  Flap suppression timer:   Off
  Cisco extensions:         Disabled
  Non-revertive:            Disabled
mLACP:
  IPv4 BFD:                 Not configured
  IPv6 BFD:                 Not configured

```

Port	Device	State	Port ID	B/W, kbps
Hu0/0/0/34	Local	Standby	0x8000, 0x0001	100000000

Link is not forwarding viable and in standby state

Configuration Examples for Link Bundling

This section contains the following examples:

Example: Configuring an Ethernet Link Bundle

The following example shows how to join two ports to form an EtherChannel bundle that runs LACP:

```

RP/0/RP0/CPU0:Router(config)# config

RP/0/RP0/CPU0:Router(config-if)# interface Bundle-Ether 3
RP/0/RP0/CPU0:Router(config-if)# ipv4 address 1.2.3.4/24
RP/0/RP0/CPU0:Router(config-if)# bundle minimum-active bandwidth 620000
RP/0/RP0/CPU0:Router(config-if)# bundle minimum-active links 1
RP/0/RP0/CPU0:Router(config-if)# exit
RP/0/RP0/CPU0:Router(config-if)# interface TenGigE 0/3/0/0
RP/0/RP0/CPU0:Router(config-if)# bundle id 3 mode active
RP/0/RP0/CPU0:Router(config-if)# no shutdown
RP/0/RP0/CPU0:Router(config)# exit
RP/0/RP0/CPU0:Router(config-if)# interface TenGigE 0/3/0/1
RP/0/RP0/CPU0:Router(config-if)# bundle id 3 mode active
RP/0/RP0/CPU0:Router(config-if)# no shutdown
RP/0/RP0/CPU0:Router(config-if)# exit

```

This example shows the configuration in the case of a mixed speed bundle:

```

RP/0/RP0/CPU0:Router(config)# config

RP/0/RP0/CPU0:Router(config-if)# interface bundle-ether 50
RP/0/RP0/CPU0:Router(config-if)# root
RP/0/RP0/CPU0:Router(config-if)# interface TenGigE 0/0/0/11
RP/0/RP0/CPU0:Router(config-if)# bundle id 50 mode active
RP/0/RP0/CPU0:Router(config-if)# no shutdown
RP/0/RP0/CPU0:Router(config-if)# interface TenGigE 0/0/0/16
RP/0/RP0/CPU0:Router(config-if)# bundle id 50 mode active
RP/0/RP0/CPU0:Router(config-if)# no shutdown
RP/0/RP0/CPU0:Router(config-if)# interface TenGigE 0/0/0/27
RP/0/RP0/CPU0:Router(config-if)# bundleid 50 mode active

```

```
RP/0/RP0/CPU0:Router(config-if)# no shutdown
RP/0/RP0/CPU0:Router(config-if)# interface HundredGigE 0/6/0/1
RP/0/RP0/CPU0:Router(config-if)# bundleid 50 mode active
RP/0/RP0/CPU0:Router(config-if)# no shutdown
RP/0/RP0/CPU0:Router(config-if)# root
RP/0/RP0/CPU0:Router(config)# commit
RP/0/RP0/CPU0:Router(config)# end
```

The following output is shown for the **show bundle bundle-ether** command:

show bundle bundle-ether50

```
Bundle-Ether50
Status: Up
Local links <active/standby/configured>: 4 / 0 / 4
Local bandwidth <effective/available>: 130000000 (130000000) kbps
MAC address (source): 0011.2233.4458 (Chassis pool)
Inter-chassis link: No
Minimum active links / bandwidth: 1 / 1 kbps
Maximum active links: 64
Wait while timer: 2000 ms
Load balancing: Default
LACP: Operational
Flap suppression timer: Off
Cisco extensions: Disabled
mLACP: Not configured
IPv4 BFD: Not configured
```

Port	Device	State	Port ID	B/W, kbps
Te0/0/0/11	Local	Active	0x8000, 0x0002	10000000
	Link is Active			
Te0/0/0/16	Local	Active	0x8000, 0x0003	10000000
	Link is Active			
Te0/0/0/27	Local	Active	0x8000, 0x0004	10000000
	Link is Active			
Hu0/6/0/1	Local	Active	0x8000, 0x0001	100000000
	Link is Active			

In order to view the weight of a mixed speed bundle, run the **show bundle load-balancing** command. The following is the truncated output of this command.

```
show bundle load-balancing bundle-ether50 location 0/0/cpu0
```

<snip>

```
Bundle-Ether50
Type: Ether (L3)
Members <current/max>: 4/64
Total Weighting: 13
Load balance: Default
Locality threshold: 65
Avoid rebalancing? False
Sub-interfaces: 1
```

```
Member Information:
Port: LON ULID BW
-----
Hu0/6/0/1 0 0 10
Te0/0/0/11 1 1 1
Te0/0/0/16 2 2 1
Te0/0/0/27 3 3 1
```

```

Platform Information:
=====

      * Bundle Summary Information *
      -----

Interface      : Bundle-Ether50   Ifhandle      : 0x00000ce0
Lag ID         : 1                Virtual Port   : 255
Number of Members : 4              Local to LC    : Yes
Hash Modulo Index : 13
MGSCP Operational Mode : No

Member Information:
LON   Interface   ifhandle   SFP  port  slot  remote/rack_id
-----
0     Hu0/6/0/1    0x100001c0 648  116  8     0/0
1     Te0/0/0/11  0x04000380 65   9    2     0/0
2     Te0/0/0/16  0x040004c0 67   8    2     0/0
3     Te0/0/0/27  0x04000780 72   4    2     0/0

</snip>

```

Example: Configuring a VLAN Link Bundle

The following example shows how to create and bring up two VLANs on an Ethernet bundle:

```

RP/0/RP0/CPU0:Router(config-subif)# config
RP/0/RP0/CPU0:Router(config-subif)# interface Bundle-Ether 1
RP/0/RP0/CPU0:Router(config-ifsubif)# ipv4 address 1.2.3.4/24
RP/0/RP0/CPU0:Router(config-ifsubif)# bundle minimum-active bandwidth 620000
RP/0/RP0/CPU0:Router(config-if)# bundle minimum-active links
RP/0/RP0/CPU0:Router(config-ifsubif)# exit
RP/0/RP0/CPU0:Router(config-subif)# ip addr 20.2.3.4/24
RP/0/RP0/CPU0:Router(config-subif)# interface Bundle-Ether 1.1
RP/0/RP0/CPU0:Router(config-subif)# encapsulation dot1q 100
RP/0/RP0/CPU0:Router(config-subif) # ip addr 10.2.3.4/24
RP/0/RP0/CPU0:Router(config-subifif)# no shutdown
RP/0/RP0/CPU0:Router(config-subifif)# exit
RP/0/RP0/CPU0:Router(config-if)# interface Bundle-Ether 1.2
RP/0/RP0/CPU0:Router(config-subif)# dot1q vlan 10
RP/0/RP0/CPU0:Router(config-subif)Router # ip addr20.2.3.4/24

RP/0/RP0/CPU0:Router(config-subifif)# no shutdown
RP/0/RP0/CPU0:Router(config-subifif)# exit
RP/0/RP0/CPU0:Router(config)# interface gig 0/1/5/7
RP/0/RP0/CPU0:Router(config-if)# bundle-id 1 mode act
RP/0/RP0/CPU0:Router(config-if)# commit
RP/0/RP0/CPU0:Router(config-if)# exit

```




CHAPTER 9

Configuring Traffic Mirroring

This module describes the configuration of the traffic mirroring feature. Traffic mirroring is sometimes called port mirroring, or switched port analyzer (SPAN).

Feature History for Traffic Mirroring

Release 7.3.1	SPAN to File feature was introduced.
Release 7.2.12	Local SPAN feature was introduced.
Release 7.0.14	Support for the following features was introduced in ERSPAN: <ul style="list-style-type: none">• Configuration of IP DSCP.• Tunnel IP.• Ability to source ranges of interfaces and SVIs.• Sequence bit is set in the GRE header and the value of sequence number is always 0 for ERSPAN packets.• ERSPAN and Security ACL should be separate. Support for File Mirroring was introduced.
Release 7.0.11	This feature was introduced.

- [Introduction to Traffic Mirroring, on page 152](#)
- [Restrictions for Traffic Mirroring, on page 161](#)
- [Configuring Traffic Mirroring, on page 163](#)
- [Attaching the Configurable Source Interface, on page 166](#)
- [Introduction to ERSPAN Rate Limit, on page 168](#)
- [Introduction to Local SPAN, on page 170](#)
- [Traffic Mirroring with DSCP, on page 174](#)
- [Monitor Multiple ERSPAN Sessions with SPAN and Security ACL, on page 179](#)
- [SPAN to File, on page 180](#)
- [Mirroring Forward-Drop Packets, on page 186](#)
- [Mirroring Buffer Drop Packets, on page 189](#)
- [Introduction to File Mirroring, on page 193](#)

- [Traffic Mirroring Configuration Examples, on page 194](#)
- [Troubleshooting Traffic Mirroring, on page 197](#)

Introduction to Traffic Mirroring


Traffic mirroring, which is sometimes called port mirroring, or Switched Port Analyzer (SPAN) is a Cisco proprietary feature. Traffic mirroring enables you to monitor Layer 3 network traffic passing in, or out of, a set of Ethernet interfaces. You can then pass this traffic to a network analyzer for analysis.

Traffic mirroring copies traffic from one or more Layer 3 interfaces or sub-interfaces. Traffic mirroring then sends the copied traffic to one or more destinations for analysis by a network analyzer or other monitoring device. Traffic mirroring does not affect the switching of traffic on the source interfaces or sub-interfaces. It allows the system to send mirrored traffic to a destination interface or sub-interface.


Traffic mirroring is introduced on switches because of a fundamental difference between switches and hubs. When a hub receives a packet on one port, the hub sends out a copy of that packet from all ports except from the one at which the hub received the packet. In case of switches, after a switch boots, it starts to build up a Layer 2 forwarding table on the basis of the source MAC address of the different packets that the switch receives. After the system builds this forwarding table, the switch forwards traffic that is destined for a MAC address directly to the corresponding port.

For example, if you want to capture Ethernet traffic that is sent by host A to host B, and both are connected to a hub, attach a traffic analyzer to this hub. All other ports see the traffic between hosts A and B.

Figure 8: Traffic Mirroring Operation on a Hub



On a switch or router, after the system learns host B MAC address, the system forwards the unicast traffic from A to B to the B port. Therefore, the traffic analyzer does not see this traffic.



In this configuration, the traffic analyzer captures only traffic that is flooded to all ports, such as:

- Broadcast traffic
- Multicast traffic with CGMP or Internet Group Management Protocol (IGMP) snooping disabled
- Unknown unicast traffic on a switch

An extra feature is necessary that artificially copies unicast packets that host A sends. This extra feature is traffic mirroring. When traffic mirroring is enabled, the traffic analyzer is attached to a port that is configured to receive a copy of every packet that host A sends. This port is called a traffic mirroring port. The other sections of this document describe how you can fine tune this feature.

From Release 24.3.1, layer 2 SPAN features are supported on PWHE. For more information about the layer 2 SPAN features supported on PWHE, see the *Traffic Mirroring on PWHE* section in the *L2VPN Configuration Guide for Cisco 8000 Series Routers*.

Implementing Traffic Mirroring on the Cisco 8000 Series Routers

ERSPAN

Table 25: Feature History Table

Feature Name	Release Information	Feature Description
Partial Packet Capture Ability for ERSPAN (RX)	Release 7.5.3	With this feature, you can perform partial packet capture in the RX direction. Earlier, the ability for entire packet capture was available, now you can choose entire or partial packet capture in the RX direction. Here, partial packet capture is also known as truncation.
ERSPAN over MPLS Traffic	Release 7.5.3	With this release, the router allows you to mirror MPLS traffic and set up the GRE tunnel with the next hop over a labeled path. This feature helps you to remote-monitor the traffic on traffic analyzers.
Higher Payload Analysis with Eight ERSPAN Sessions	Release 7.3.2	With this release, Cisco 8000 Series routers support eight ERSPAN sessions. This functionality helps you analyze higher payloads in real time across Layer 3 domains on your network.

Feature Name	Release Information	Feature Description
ERSPAN over GRE IPv6	Release 7.3.2	With this release, the router allows you to mirror IPv4 or IPv6 traffic with ERSPAN over GRE IPv6 sessions to monitor traffic on remote traffic analyzers. In earlier releases, ERSPAN traffic monitoring was possible only on IPv4 networks.

Encapsulated Remote Switched Port Analyzer (ERSPAN) mirrors traffic on one or more source ports and delivers the mirrored traffic to destination port on another switch or management server. ERSPAN enables network operators to troubleshoot issues in the network in real-time using automated tools that auto-configures ERSPAN parameters on the network devices to send specific flows to management servers for in-depth analysis.


ERSPAN transports mirrored traffic over an IP network. The traffic is encapsulated at the source router and is transferred across the network.

From Cisco IOS XR Software Release 7.5.3 onwards, the packet truncation feature is supported over remote GRE tunnels. You can now get the flexibility to truncate packets and mirror the traffic.

Starting with Cisco IOS XR Software Release 7.0.14, sequence bit is set in the GRE header and the value of sequence number is always 0 for ERSPAN packets.

Starting with Cisco IOS XR Software Release 7.5.3, the sequence number bit will always be set to one and the sequence number field (4 bytes), will always be set to zero.

Figure 9: ERSPAN over GRE



Supported Capabilities

The following capabilities are supported:

- The source interfaces are layer 3 interfaces, such as physical, and bundle interfaces or subinterface.
- The routers mirror IPv4 and IPv6 traffic.
- ERSPAN with GRE IPv4 or IPv6 has tunnel destinations.

- ERSPAN supports only RX direction.
- ERSPAN over GRE IPv4 and IPv6 supports SPAN ACL.
- Supports MPLS traffic mirroring and GRE tunnel configuration with the next hop over a labeled path.
- Each monitor session allows only one destination interface.
- ACL permit or deny entries with capture action are part of mirroring features.
- The next hop interface must be a main interface. It can be a Physical or Bundle interface.
- Supports full packet capture.
- In ERSPAN over GRE IPv6, the **HopLimit** and **TrafficClass** fields in outer IPv6 header are editable under the tunnel configuration.
- The maximum SPAN sessions supported in the Cisco 8000 Router are as follows:.

SPAN Type	7.3.1 and Prior Releases	7.3.2 and Later Releases
ERSPAN (GRE IPv4, GRE IPv6, or GRE IPv4 + GRE IPv6)	4	8
Local SPAN	4	4
SPAN to File	4	4
Combined SPAN (GRE IPv4 + GRE IPv6 + Local SPAN + SPAN to File)	4	8

- Starting with Release 24.2.11, on all Egress Traffic Management (ETM)-based platforms, when the NPU compatibility mode is set to P100, the maximum number of SPAN sessions supported on the 88-LC1-52Y8H-EM and 88-LC1-12TH24FH-E line cards are as follows:
 - ERSPAN (GRE IPv4, GRE IPv6, or GRE IPv4 + GRE IPv6): 4
 - Local SPAN: 4
 - SPAN to File: 4
 - Combined SPAN (GRE IPv4 + GRE IPv6 + Local SPAN + SPAN to File): 4



Note For more information on NPU compatibility mode, see [Configure the Comaptibility Mode](#).

Supported Capabilities for ERSPAN Packet Truncation support

The following are the capabilities and requirements:

- Ability to enable the new ERSPAN GREv4 and GREv6 truncation configuration per device.
- Truncation configuration should be on the monitor sessions. Packets received from all sources will only be truncated when you configure the truncation on a monitor session.

- By default, the whole packet will be mirrored without the **mirror first <number>** (truncation size) configuration.
- If the monitor session truncation size is less than the configured-truncation size (343 bytes), then whole packet is mirrored.

If the monitor session truncation size exceeds 343 bytes, the configuration is accepted. However, only 343 bytes truncation size is programmed.

An `ios-msg` is displayed to warn the user.

Example: ERSPAN only support 343 bytes truncation size. `monitor-session` with `session_id <id>` will be set to 343 bytes only.

Restrictions

The following are the ERSPAN and SPAN ACL restrictions:

- The ERSPAN mirror packet is received with a TTL minus 1.
The mirror packet is not identical to the incoming packet and TTL minus 1 is the expected value in the ERSPAN packet.
- The router mirrors only unicast traffic.
However, from Cisco IOS XR Software Release 7.5.3 onwards, the router can mirror multicast traffic.
- Remove and re-apply monitor-sessions on all interfaces after modifying the access control list (ACL).
- GRE tunnel is only dedicated to ERSPAN mirrored packets. There should be no IPv4 and IPv6 address configured under the GRE tunnel.
- Only ERSPAN TYPE II header is supported. The value of the index field is always 0. The value of the session-ID field is an internal number that is used by the data path to distinguish between sessions.
- Traffic accounting of the ERSPAN mirrored packets is not supported.



Note You can view the SPAN packet count per session, using the [show monitor-session status internal](#) command.

- ERSPAN decapsulation is unsupported.
- From Cisco IOS XR Software Release 7.5.3 onwards, the ERSPAN will be functional regardless of any configuration related to MPLS or LDP present on the router.
- MPLS packet mirroring is supported only from Cisco IOS XR Software Release 7.5.3 onwards.
- Due to data path limitation, the source IPv6 addresses of the outer IPv6 header of the ERSPAN packet have only higher 64 bits as valid. The lower 64-bits value is changed to zero. The destination GREv6 IPv6 address should contain all the 128 bits.

Traffic Mirroring Terminology

- Ingress traffic—Traffic that enters the switch.

- Egress traffic—Traffic that leaves the switch.
- Source port—A port that the system monitors with the use of traffic mirroring. It is also called a monitored port.
- Destination port—A port that monitors source ports, usually where a network analyzer is connected. It is also called a monitoring port.
- Monitor session—A designation for a collection of traffic mirroring configurations consisting of a single destination and, potentially, many source interfaces.

Characteristics of the Source Port

A source port, also called a monitored port, is a switched or routed port that you monitor for network traffic analysis. In a single local or remote traffic mirroring session, you can monitor source port traffic, such as received (Rx) for ingress traffic. Your router can support any number of source ports (up to a maximum number of 800).

A source port has these characteristics:


- It can be any port type, such as Bundle Interface, sub-interface, 100-Gigabit Ethernet, or 400-Gigabit Ethernet.



Note Bridge group virtual interfaces (BVI) are not supported.

- Each source port can be monitored in only one traffic mirroring session.
- It cannot be a destination port.
- Each source port can be configured with a direction (ingress) to monitor. For bundles, the monitored direction applies to all physical ports in the group.

Figure 10: Network Analysis on a Router with Traffic Mirroring



In the figure above, the network analyzer is attached to a port that is configured to receive a copy of every packet that host A sends. This port is called a traffic mirroring port.

Characteristics of the Monitor Session

A monitor session is a collection of traffic mirroring configurations consisting of a single destination and, potentially, many source interfaces. For any given monitor session, the traffic from the source interfaces

(called *source ports*) is sent to the monitoring port or destination port. If there is more than one source port in a monitoring session, the traffic from the several mirrored traffic streams is combined at the destination port. The result is that the traffic that comes out of the destination port is a combination of the traffic from one or more source ports.

Monitor sessions have these characteristics:

- A single monitor session can have only one destination port.
- A single destination port can belong to only one monitor session.



Note The destination of ERSPAN monitoring session is a GRE IPv4 or IPv6 tunnel.

Supported Traffic Mirroring Types

The system supports the following traffic mirroring types:

- ACL-based traffic mirroring. The system mirrors traffic that is based on the configuration of the global interface ACL.
- Layer 3 traffic mirroring is supported. The system can mirror Layer 3 source ports.

ACL-Based Traffic Mirroring

You can mirror traffic that is based on the definition of a global interface access list (ACL). When you are mirroring Layer 3 traffic, the ACL is configured using the **ipv4 access-list** or **ipv6 access-list** command with the **capture** keyword. The **permit** and **deny** commands determine the behavior of regular traffic. The **capture** keyword designates that the packet is to be mirrored to the destination port.

Starting with Cisco IOS XR Software Release 7.0.14, configuration of ERSPAN and security ACL will be separate. Neither of these will have an impact or dependency on the other, but both can be applied simultaneously.

ERSPAN over GRE IPv6

The ERSPAN over GRE IPv6 feature enables mirroring IPv4 or IPv6 traffic in your network. The router encapsulates the traffic adding an ERSPAN header inside the GRE IPv6 packet. The GRE header of the ERSPAN encapsulated packets have the sequence number set to 0. The router sends the replicated traffic packet to be monitored to the destination through the GRE IPv6 channel to achieve traffic mirroring. The mirrored traffic is sent to remote traffic analyzer for monitoring purposes. For the traffic mirroring to work, the ERSPAN GRE IPv6 tunnel next-hop must have ARP or neighbor resolved. We recommend using the `cef proactive-arp-nd enable` command to configure missing adjacency information for the next hop.



Note The GRE tunnel configured for ERSPAN should only be used for mirrored traffic. There should be no IPv4 or IPv6 address configured under the GRE Tunnel.

```
Router# configure
Router(config)# cef proactive-arp-nd enable
Router(config)# commit
```


Configuring ERSpan over GRE IPv6

1. Enable GRE IPv6 tunnel configuration.

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#interface tunnel-ip1
RP/0/RP0/CPU0:router(config-if)#tunnel mode gre ipv6
RP/0/RP0/CPU0:router(config-if)#tunnel source 2001:DB8:1::1
RP/0/RP0/CPU0:router(config-if)#tunnel destination 2001:DB8:2::1
RP/0/RP0/CPU0:router(config-if)#no shut
RP/0/RP0/CPU0:router(config)#commit
```

2. Enable ERSpan session.

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#monitor-session mon1 ethernet
RP/0/RP0/CPU0:router(config-mon)#destination interface tunnel-ip1
RP/0/RP0/CPU0:router(config-mon)#commit
RP/0/RP0/CPU0:router(config-mon)#end
```

3. Configure ERSpan session under port to be monitored.

```
RP/0/RP0/CPU0:router(config)#interface HundredGigE0/1/0/14
RP/0/RP0/CPU0:router(config-if)#monitor-session mon1 ethernet direction rx-only
RP/0/RP0/CPU0:router(config-if-mon)#exit
RP/0/RP0/CPU0:router(config-if)#exit
RP/0/RP0/CPU0:router(config)#interface Bundle-Ether1
RP/0/RP0/CPU0:router(config-if)#monitor-session mon1 ethernet direction rx-only
RP/0/RP0/CPU0:router(config-if-mon)#exit
RP/0/RP0/CPU0:router(config-if)#exit
RP/0/RP0/CPU0:router(config)#interface HundredGigE0/1/0/15.100
RP/0/RP0/CPU0:router(config-subif)#monitor-session mon1 ethernet direction rx-only
```

Verification

Use the `show monitor-session status` command to verify the configuration of the ERSpan over GRE IPv6 feature.

```
P/0/RP0/CPU0:router#show monitor-session mon1 status
Monitor-session mon1
Destination interface tunnel-ip1
```

```
=====
Source Interface          Dir      Status
-----
Hu0/1/0/14                Rx      Operational
Hu0/1/0/15.100           Rx      Operational
BE1                       Rx      Operational
BE1.1                     Rx      Operational
```

```
RP/0/RP0/CPU0:R1-SF-D#show monitor-session erspan3 status internal
```

```
Thu Jul 15 06:00:14.720 UTC
Information from SPAN Manager and MA on all nodes:
Monitor-session erspan3 (ID 0x00000007) (Ethernet)
SPAN Mgr: Destination interface tunnel-ip372 (0x0f00049c)
Last error: Success
Tunnel data:
  Mode: GREoIPv6
  Source IP: 77:3:1::79
  Dest IP: 95::90
  VRF:
  ToS: 100
  TTL: 200
  DFbit: Not set
0/3/CPU0: Destination interface tunnel-ip372 (0x0f00049c)
```

```

Tunnel data:
  Mode: GREoIPv6
  Source IP: 77:3:1::79
  Dest IP: 95::90
  VRF:
  ToS: 100
  TTL: 200
  DFbit: Not set
0/RP0/CPU0: Destination interface tunnel-ip372 (0x0f00049c)
Tunnel data:
  Mode: GREoIPv6
  Source IP: 77:3:1::79
  Dest IP: 95::90
  VRF:
  ToS: 100
  TTL: 200
  DFbit: Not set
Information from SPAN EA on all nodes:
Monitor-session 0x00000007 (Ethernet)
0/3/CPU0: Name 'erspan3', destination interface tunnel-ip372 (0x0f00049c)
Platform, 0/3/CPU0:
  Monitor Session ID: 7

  Monitor Session Packets: 2427313444
  Monitor Session Bytes: 480591627492

```

Configuring Partial Packet Capture Ability for ERSPAN (RX)

To configure partial traffic mirroring, use the **mirror first** command in monitor session configuration mode.

Mirror first <number>: Configures the size of truncation packets for an ERSPAN session

Use the following command to create a ERSPAN monitor session for mirroring the packets:

```

monitor-session <name> [ethernet]
destination interface tunnel-ip <number>
mirror first <number>
  traffic-class <traffic-class>

```

Configuration Example

Use the following command to create a ERSPAN monitor session for mirroring packets to Tunnel-IP 30 with truncation enabled:

```

monitor-session mon1 ethernet
  destination interface tunnel-ip 30
  mirror first 343
!

```

Attach the session to the interfaces using the following configuration:

```

interface <>
  monitor-session session-name ethernet direction rx-only|tx-only|both | acl [acl_name]

```

Running Configuration

```

interface tunnel-ip30
  tunnel mode gre ipv4
  tunnel source 2.2.2.2
  tunnel destination 200.0.0.2
!

interface HundredGigE0/0/0/12
  ipv4 address 12.0.0.2 255.255.255.0
  monitor-session mon1 ethernet direction rx-only
!

```

Verification

The **show monitor-session status internal** displays the size of the programmed truncation.

Example:

```
Router#show monitor-session mon1 status internal
Fri Apr 12 18:50:45.006 UTC
Information from SPAN Manager and MA on all nodes:
Packet truncation size: 343B
Monitor-session mon1 (ID 0x00000001) (Ethernet)
SPAN Mgr: Destination interface Tunnel-IP 20 (0x0f000250)
          Last error: Success

Information from SPAN EA on all nodes:
Monitor-session 0x00000001 (Ethernet)
O/RP0/CPU0: Name 'mon1', destination interface Tunnel-IP 20 (0x0f000250)
Platform, O/RP0/CPU0:

Monitor Session Packets: 142462

Monitor Session Bytes: 7653237
```

ERSPAN Traffic to a Destination in a Non-Default VRF

Table 26: Feature History Table

Feature Name	Release Information	Description
ERSPAN Traffic to a Destination in a Non-Default VRF	Release 7.5.2 Release 7.3.4	Encapsulated Remote Switched Port Analyzer (ERSPAN) now transports mirrored traffic through GRE tunnels with multiple VRFs, helping you design your network with multiple Layer 3 partitions. In earlier releases, ERSPAN transported mirrored traffic through GRE tunnels that belonged to only default VRF.

Restrictions for Traffic Mirroring

The system supports the following forms of traffic mirroring:

- Mirroring traffic to a GRE IPv4 or IPv6 tunnel (also known as Encapsulated Remote Switched Port Analyzer [ER-SPAN] in Cisco IOS Software). The system allows 8 monitor sessions for ERSPAN, 4 monitor sessions for Local SPAN, and 4 monitor sessions for SPAN to File. The total number of monitor sessions for all SPAN features is 8.
- The system does not support traffic mirroring counters per interface.
- The system does not support bundle member interfaces as sources for mirroring sessions.
- ERSPAN tunnel statistics is not supported.

The following general restrictions apply to traffic mirroring using ACLs:

- Configure ACLs on the source interface to avoid default mirroring of traffic. If a Bundle interface is a source interface, configure the ACLs on the bundle interface (not bundle members).

The following restrictions apply to ERSPAN ACL:

- ERSPAN next-hop must have ARP resolved.
 - Any other traffic or protocol triggers ARP.
- ERSPAN decapsulation is not supported.
- ERSPAN does not work if the GRE next hop is reachable over subinterface. For ERSPAN to work, the next hop must be reachable over the main interface.
- However, from Cisco IOS XR Software Release 7.5.3 onwards, GRE next hop can be resolved over subinterface or the main interface.

Modifying ERSPAN monitor-session configuration

When you modify the ERSPAN monitor-session configuration, the **show configuration** and **show configuration commit changes** command outputs differ. Specifically, the **show configuration commit changes** command output displays some extraneous ACL commands deleted and added back. This modified output doesn't impact your configuration or affect performance. This issue is fixed in Cisco IOS XR Release 7.5.1.

The following example highlights the extraneous ACL commands under the **show configuration commit changes** command output.

```
Router(config)#interface HundredGigE0/1/0/0
Router(config-if)#no monitor-session ERSPANtun2005
Router(config-if)#monitor-session ERSPANtun2 ethernet direction rx-only port-level
Router(config-if-mon)#acl
Router(config-if-mon)#acl ipv4 erspan-filter
Router(config-if-mon)#acl ipv6 erspan-filter-ipv6
Router(config-if-mon)#
Router(config-if-mon)#show configuration
Building configuration...
!! interface HundredGigE0/1/0/0
   monitor-session ERSPANtun2 ethernet direction rx-only port-level
      acl
      acl ipv4 erspan-filter
      acl ipv6 erspan-filter-ipv6
   !
!
end

Router(config-if-mon)#commit
Router(config-if-mon)#end
Router#sh configuration commit changes las 1
Building configuration...
!!
interface HundredGigE0/1/0/0
no monitor-session ERSPANtun2005 ethernet direction rx-only port-level
monitor-session ERSPANtun2 ethernet direction rx-only port-level
no acl
acl
no acl ipv4 erspan-filter
acl ipv4 erspan-filter
```

```

no acl ipv6 erspan-filter-ipv6
acl ipv6 erspan-filter-ipv6
!
!
end

```

Configuring Traffic Mirroring

These tasks describe how to configure traffic mirroring:

Configuring ACLs for Traffic Mirroring

This section describes the configuration for creating ACLs for traffic mirroring. You must configure the global interface ACLs by using one of the following commands with the **capture** keyword:

- **ipv4 access-list**
- **ipv6 access-list**



Note Starting with Cisco IOS XR Software Release 7.0.14, ACL feature will provide a support of separate ACL configuration for SPAN.

Configuration

• Security ACL

Use the following configuration to configure ACLs for traffic mirroring.

```

/* Create an IPv4 ACL (TM-ACL) for traffic mirroring */
Router(config)# ipv4 access-list TM-ACL
Router(config-ipv4-acl)# 10 permit udp 10.10.10.0 0.0.0.255 eq 10 any capture
Router(config-ipv4-acl)# 20 permit udp 10.10.10.0 0.0.0.255 eq 20 any
Router(config-ipv4-acl)# exit
Router(config)# commit

/* Apply the traffic monitoring to SPAN source interface */
Router(config)# interface HundredGigE0/0/0/12
Router(config-if)# monitor-session mon1 ethernet direction rx-only port-level acl
Router(config-if)# ipv4 access-group TM-ACL ingress
!

```

Use the following configuration as an example to deny data forwarding for an ACE entry, but still mirror the traffic:

```

ipv4 access-list acl1
10 deny ipv4 any 2.1.0.0/16 capture
20 permit ipv4 any any
!

```

If **acl1** is attached to the interface as shown below:

```

RP/0/RP0/CPU0(config-if)# ipv4 access-group acl1 ingress

```

Data Traffic to 2.1.0.0/16 is dropped. Mirroring happens only if `icmp-off` keyword is added to the ACE as shown below. If this keyword is not added, mirroring does not take place. Furthermore, the `icmp-off` workaround is applicable only to security ACL.

```
ipv4 access-list acl1
10 deny ipv4 any 2.1.0.0/16 capture icmp-off
20 permit ipv4 any any
!
```

• SPAN ACL

- SPAN ACL does not support User Defined Fields (UDF).
- Deny action in SPAN ACL is ignored, and no packet drops from SPAN ACL. Deny ACEs will be internally converted to permit ACEs. Packets will also be mirrored.
- There is no implicit deny-all entry in SPAN ACL.
- IPV6 ACL is required for mirroring IPV6 packet, if IPV4 ACL is configured, and vice versa. This follows the same structure as Security ACL with IPv4 and IPv6 mirror options.

Use the following configuration to enable traffic mirroring with ACLs.

```
/* Create a SPAN IPv4 ACL (v4-monitor-acl) for traffic mirroring */
Router(config)# ipv4 access-list v4-monitor-acl
Router(config-ipv4-acl)# 10 permit udp 20.1.1.0 0.0.0.255 eq 10 any
Router(config-ipv4-acl)# 20 permit udp 30.1.1.0 0.0.0.255 eq 20 any
Router(config-ipv4-acl)# exit
Router(config)# commit

/*Create a SPAN IPv6 ACL (v6-monitor-acl) for traffic mirroring */
Router(config)# ipv6 access-list v6-monitor-acl
Router(config-ipv6-acl)# 10 permit ipv6 host 120:1:1::1 host 130:1:1::1
Router(config-ipv6-acl)# exit

/* Apply the traffic monitoring to SPAN source interface */
Router(config)# interface HundredGigE0/0/0/12
Router(config-if)# monitor-session mon1 ethernet direction rx-only
Router(config-if)# acl ipv4 v4-monitor-acl
Router(config-if)# acl ipv6 v6-monitor-acl!
```



Note For SPAN to work, the `capture` keyword is required for Security ACL.

Use the `show access-lists [ipv4 | ipv6] acl-name hardware ingress span [detail | interface | location | sequence | verify] location x command` to display ACL information:

```
Router# show access-lists ipv4 v4span1 hardware ingress span interface bundle-Ether 100
location 0/3/cpu0
ipv4 access-list v4span1
10 permit ipv4 host 51.0.0.0 host 101.0.0.0
20 permit ipv4 host 51.0.0.1 host 101.0.0.1
30 permit ipv4 host 51.0.0.2 any
40 permit ipv4 any host 101.0.0.3
50 permit ipv4 51.0.1.0 0.0.0.255 101.0.1.0 0.0.0.255
60 permit ipv4 51.0.2.0 0.0.0.255 101.0.2.0 0.0.0.255 precedence critical
```

Troubleshooting ACL-Based Traffic Mirroring

Take note of these configuration issues:

- Even when the system configures the **acl** command on the source mirroring port, if the ACL configuration command does not use the **capture** keyword, the system does not mirror traffic.
- If the ACL configuration uses the **capture** keyword, but you have not configured the **acl** command on the source port, the system mirrors the traffic, but does not apply access list configuration.

This example shows both the **capture** keyword in the ACL definition and the **acl** command that is configured on the interface:

```
/* Create an IPv4 ACL (TM-ACL) for traffic mirroring */
Router(config)# ipv4 access-list TM-ACL
Router(config-ipv4-acl)# 10 permit udp 10.1.1.0 0.0.0.255 eq 10 any capture
Router(config-ipv4-acl)# 20 permit udp 10.1.1.0 0.0.0.255 eq 20 any

Apply the traffic monitoring to interface
Router(config)#interface HundredGigE0/0/0/12
Router(config-if)# monitor-session mon1 ethernet direction rx-only port-only acl
Router(config-if)# ipv4 access-group TM-ACL ingress
```

Flexible CLI for ERSPAN

Starting with Cisco IOS XR Software Release 7.0.14, ERSPAN can be configured using flexible CLI. This CLI is a single configuration object containing all the properties of an ERSPAN session, tunnel properties, and the list of source interfaces, which can be easily removed and re-added. Flexible CLI minimises risk of user error and promotes operational simplicity.

Configure a flexible CLI group in ERSPAN containing:

- Global ERSPAN session configuration
- Tunnel interface configuration
- ERSPAN source attachment configuration, applied to a regexp of interface names



Note The flexible CLI group contains only the session and interface properties. The session and interface objects themselves must be created in the configuration as usual.

The following example shows a global flexible CLI configuration:

```
group erspan-group-foo
  monitor-session 'foo' ethernet /* Global configuration */
  destination interface tunnel-ip0
  !
  interface 'tunnel-ip0' /* Tunnel interface configuration */
  tunnel tos 10
  tunnel mode gre ipv4
  tunnel source 10.10.10.1
  tunnel destination 20.20.20.2
  !
  interface 'GigabitEthernet0/0/0/[0-3]' /* Interface configuration */
  monitor-session foo ethernet
```

```
!
end-group
```

To enable all ERSPAN configurations, execute `apply-group erspan-group-foo` command. To disable ERSPAN configuration, delete this command.



Note The following three keywords are regular expressions and must be quoted:

- Definition of session name (example: `foo`)
- Definition of tunnel name (example: `tunnel-ip0`)
- Set of source interface names (example: `GigabitEthernet0/0/0/[0-3]`)

Use the `show running-config inheritance` command to view the final configuration after the group is expanded, and the `show monitor-session status` to check the operational state of ERSPAN session.



Note Starting from Release 7.3.3, when a combination of IP-in-IP decap and GRE ERSPAN tunnels are in use, resource utilization of IP-in-IP decap tunnels is accounted. However, resource utilization of ERSPAN GRE tunnels is not accounted in the *Total In Use* counter of **show controllers npu resources sipidxtbl location all** command output, but the *OOR State* would display *RED* if the total number of IP-in-IP decap and ERSPAN GRE tunnels reach 15.

Attaching the Configurable Source Interface

Procedure

Step 1 **configure**

Example:

```
RP/0/RP0/CPU0# configure
```

Enters global configuration mode.

Step 2 **interface type number**

Example:

```
RP/0/RP0/CPU0(config)# interface HundredGigE 0/1/0/10/0/1/0
```

Enters interface configuration mode for the specified source interface. The interface number is entered in *rack/slot/module/port* notation. For more information about the syntax for the router, use the question mark (?) online help function.

Step 3 **ipv4 access-group acl-name {ingress | egress}**

Example:


```
RP/0/RP0/CPU0(config-if)# ipv4 access-group acl1 ingress
```

Controls access to an interface.

Step 4 **monitor-session** *session-name* **ethernet direction rx-only**port-level

Example:

```
RP/0/RP0/CPU0(config-if)# monitor-session mon1 ethernet direction rx-only port-level acl
RP/0/RP0/CPU0(config-if-mon)#
```

Attaches a monitor session to the source interface and enters monitor session configuration mode.

Note **rx-only** specifies that only ingress traffic is replicated.

Step 5 **acl**

Example:

```
RP/0/RP0/CPU0(config-if-mon)# acl
```

Specifies that the traffic mirrored is according to the defined ACL.

Note If an ACL is configured by name then this overrides any ACL that may be configured on the interface.

Step 6 **exit**

Example:

```
RP/0/RP0/CPU0(config-if-mon)# exit
RP/0/RP0/CPU0(config-if)#
```

Exits monitor session configuration mode and returns to interface configuration mode.

Step 7 **end** or **commit**

Example:

```
RP/0/RP0/CPU0(config-if)# end
```

or

```
RP/0/RP0/CPU0(config-if)# commit
```

Saves configuration changes.

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting (yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.

- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.

- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.

- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Step 8 `show monitor-session [session-name] status [detail] [error]`

Example:

```
RP/0/RP0/CPU0# show monitor-session status
```

Displays information about the monitor session.

Introduction to ERSPAN Rate Limit

With ERSPAN rate limit feature, you can monitor traffic flow through any IP network. This includes third-party switches and routers.

ERSPAN operates in the following modes:

- ERSPAN Source Session – box where the traffic originates (is SPANned).
- ERSPAN Termination Session or Destination Session – box where the traffic is analyzed.

This feature provides rate limiting of the mirroring traffic. With rate limiting, you can limit the amount of traffic to a specific rate, which prevents the network and remote ERSPAN destination traffic overloading. Be informed, if the rate-limit exceeds then the system may cap or drop the monitored traffic.

You can configure the QoS parameters on the traffic monitor session.


- Traffic Class (0 through 7)
 - Traffic class 0 has the lowest priority and 7 the highest.
 - The default traffic class is the same as that of the original traffic class.

Benefits

With ERSPAN rate limit feature, you can limit the mirrored traffic and use the mirrored traffic for data analysis.

Topology

Figure 11: Topology for ERSPAN Rate Limit



The encapsulated packet for ERSPAN is in ARPA/IP format with GRE encapsulation. The system sends the GRE tunneled packet to the destination box identified by an IP address. At the destination box, SPAN-ASIC

decodes this packet and sends out the packets through a port. ERSPAN rate limit feature is applied on the router interface to rate limit the monitored traffic.

The intermediate switches carrying ERSPAN traffic from source session to termination session can belong to any L3 network.

Configure ERSPAN Rate Limit

Use the following steps to configure ERSPAN rate limit:

```
monitor-session ERSPAN ethernet
destination interface tunnel-ip1
!

RP/0/RP0/CPU0:pyke-008#sh run int tunnel-ip 1

interface tunnel-ip1
ipv4 address 4.4.4.1 255.255.255.0
tunnel mode gre ipv4
tunnel source 20.1.1.1
tunnel destination 20.1.1.2
!

RP/0/RP0/CPU0:pyke-008#sh run int hundredGigE 0/0/0/16

interface HundredGigE0/0/0/16
ipv4 address 215.1.1.1 255.255.255.0
ipv6 address 3001::2/64
monitor-session ERSPAN ethernet direction rx-only port-level
  acl
!
ipv4 access-group ACL6 ingress
```

Running Configuration

```
!!A traffic class needs to be configured under the monitor session.
monitor-session mon2 ethernet
destination interface tunnel-ip30
traffic class 5
```

A shaper needs to be configured for this traffic class:

```
policy-map m8
class TC1
  bandwidth percent 11
!
class TC2
  bandwidth percent 12
!
class TC3
  bandwidth percent 13
!
class TC4
  bandwidth percent 14
!
class TC5
  shape average percent 15
!
class TC6
  bandwidth percent 16
!
class TC7
  bandwidth percent 17
```

```
This policy-map has to be installed on the interface over which the mirrored traffic is
sent in the egress direction:
interface TenGigE0/6/0/9/0
service-policy output m8
```

Verification

```
RP/0/RP0/CPU0:ios#show monitor-session FOO status detail
Wed May 2 15:14:05.762 UTC
Monitor-session FOO
Destination interface tunnel-ip100
Source Interfaces
-----
TenGigE0/6/0/4/0
Direction: Both
Port level: True
ACL match: Disabled
```

Introduction to Local SPAN

Local SPAN Overview

Local SPAN is the most basic form of traffic mirroring. In Local SPAN, both mirror source and mirror destination interfaces are present on the same router.

Local SPAN Supported Capabilities

The following capabilities are supported for Local SPAN:

- Only ingress traffic.
- The destination interface can only be an L2 or L3 physical main interface.
- The following interfaces are configured as sources for a Local SPAN session:
 - L3 physical main and sub-interface and bundle main and sub-interface.
 - L2 ethernet interfaces: Ethernet Flow Point (EFP) and trunk
 - BVI interface
- The following types of traffic are mirrored Local SPAN:
 - IPv4, IPv6, and MPLS
 - IP-in-IP
- Extended ACL to reduce mirrored traffic throughput
- Traffic shaping on the destination interface
- Session statistics. There's one counter for all types of traffic, that is, IPv4, IPv6, and MPLS.
- Up to four Local SPAN sessions. This session number is shared between ERSPAN, Local SPAN, and SPAN to File features.

- Up to 1000 source interfaces

Local SPAN Restrictions

The following are the restrictions for Local SPAN:

- Egress mirroring isn't supported
- The physical interface used as destination can't be a bundle member link
- GRE tunnel isn't supported as source interface and destination interface
- Port-level monitoring isn't supported
- Per-source interface mirroring statistics isn't supported. However, SPAN session statistics are supported. The session statistics would contain total number of packets mirrored by the session.
- A destination interface can't be a mirrored source interface and vice versa.
- ACL for Local SPAN is only applied in ingress direction.
- If ACL keyword is present in monitor-session configuration for an interface but no ACL is applied to that interface, traffic packets won't be mirrored
- No ACL support for MPLS traffic
- No support for NetFlow or sFlow configuration on the same interface which has Local SPAN session already configured.

Configuring Local SPAN

Configuring Local SPAN consists of 2 parts:

1. Creating a local SPAN session

```
RP/0/RP0/CPU0:router#configure
RP/0/RP0/CPU0:router(config)#monitor-session mon1 ethernet
RP/0/RP0/CPU0:router(config-mon)#destination interface HundredGigE0/1/0/0
RP/0/RP0/CPU0:router(config-mon)#commit
RP/0/RP0/CPU0:router(config-mon)#end
RP/0/RP0/CPU0:router#
```

2. Attaching the SPAN session to an interface

```
RP/0/RP0/CPU0:router(config-mon)#interface HundredGigE0/1/0/2
RP/0/RP0/CPU0:router(config-if)# monitor-session mon1 ethernet direction rx-only
RP/0/RP0/CPU0:router(config-if-mon)# no shut
RP/0/RP0/CPU0:router(config-if)#!
RP/0/RP0/CPU0:router(config-if)#
RP/0/RP0/CPU0:router(config-if)#interface Bundle-Ether1
RP/0/RP0/CPU0:router(config-if)# monitor-session mon1 ethernet direction rx-only
RP/0/RP0/CPU0:router(config-if-mon)# no shutdown
RP/0/RP0/CPU0:router(config-if)#!
RP/0/RP0/CPU0:router(config-if)#

RP/0/RP0/CPU0:monitor(config-if)#
RP/0/RP0/CPU0:monitor(config-if)#interface HundredGigE0/1/0/14.100
RP/0/RP0/CPU0:monitor(config-subif)# monitor-session mon1 ethernet direction rx-only
RP/0/RP0/CPU0:monitor(config-if-mon)# no shut
RP/0/RP0/CPU0:monitor(config-subif)#!
```

```

RP/0/RP0/CPU0:monitor(config-subif)#
RP/0/RP0/CPU0:monitor(config-subif)#interface Bundle-Ether1.1
RP/0/RP0/CPU0:monitor(config-subif)# monitor-session mon1 ethernet direction rx-only
RP/0/RP0/CPU0:monitor(config-if-mon)# no shut
RP/0/RP0/CPU0:monitor(config-subif)#!
RP/0/RP0/CPU0:monitor(config-subif)#commit

```

Verification

```

RP/0/RP0/CPU0:router#show monitor-session status
Monitor-session mon1
Destination interface HundredGigE0/1/0/0

```

```

=====
Source Interface      Dir      Status
-----
Hu0/1/0/2            Rx      Operational
Hu0/1/0/14.100      Rx      Operational
BE1                  Rx      Operational
BE1.1                Rx      Operational

```

Execute the `show monitor-session status internal` command for session statistics:

```

RP/0/RP0/CPU0:router#show monitor-session status internal
Thu Aug 13 20:05:23.478 UTC
Information from SPAN Manager and MA on all nodes:
Monitor-session mon1 (ID 0x00000001) (Ethernet)
SPAN Mgr: Destination interface HundredGigE0/1/0/0 (0x00800190)
          Last error: Success
0/1/CPU0: Destination interface HundredGigE0/1/0/0 (0x00800190)
0/RP0/CPU0: Destination interface HundredGigE0/1/0/0 (0x00800190)
Information from SPAN EA on all nodes:
Monitor-session 0x00000001 (Ethernet)
0/1/CPU0: Name 'mon1', destination interface HundredGigE0/1/0/0 (0x00800190)
Platform, 0/1/CPU0:

Monitor Session ID: 1
Monitor Session Packets: 32
Monitor Session Bytes: 4024

0/2/CPU0: Name 'mon1', destination interface HundredGigE0/1/0/0 (0x00800190)
Platform, 0/2/CPU0:

Monitor Session ID: 1
Monitor Session Packets: 0
Monitor Session Bytes: 0

```

Local SPAN with ACL

Local SPAN with ACL is used to filter and mirror ingress traffic. Only Access Control Entries (ACEs) with `capture` keyword are considered for mirroring. Both permit and deny packets are captured if the ACE contains `capture` keyword. Per interface, only one IPv4 ingress ACL and one IPv6 ingress ACL is allowed.

Configuring Local SPAN with ACL

Use the following configuration to enable local SPAN with IPv4 ACLs:

1. Configure ACLs for traffic mirroring.

```

Router(config)# ipv4 access-list acl1
Router(config-ipv4-acl)# 10 permit ipv4 25.0.0.0 0.0.0.255 any capture
Router(config-ipv4-acl)# 20 permit ipv4 20.0.0.0 0.0.0.255 any

```

```
Router(config-ipv4-acl)# 30 permit ipv4 131.1.1.0 0.0.0.255 any capture
Router(config-ipv4-acl)# 40 permit ipv4 191.1.1.0 0.0.0.255 any capture
```

2. Apply the traffic monitoring to an interface.

```
Router(config)# interface HundredGigE0/1/0/2
Router(config-if)# ipv4 address 131.1.1.2 255.255.255.0
Router(config-if)# monitor-session mon1 ethernet direction rx-only port-level
Router(config-if-mon)# acl
Router(config-if-mon)# ipv4 access-group acl1 ingress
```

Verification

```
RP/0/RP0/CPU0:ios#show running-config ipv4 access-list acl1
Thu Aug 13 20:22:54.388 UTC
ipv4 access-list acl1
 10 permit ipv4 22.0.0.0 0.0.0.255 any capture
 20 permit ipv4 20.0.0.0 0.0.0.255 any
 30 permit ipv4 131.1.1.0 0.0.0.255 any capture
 40 deny ipv4 181.1.1.0 0.0.0.255 any capture
!
```

Use the following configuration to enable local SPAN with IPv6 ACLs:

1. Configure ACLs for traffic mirroring.

```
Router(config)# ipv6 access-list acl2
Router(config-ipv6-acl)# 10 permit ipv6 10:1:1::2/64 any capture
Router(config-ipv6-acl)# 20 permit ipv6 10:1:1::3/64 any
Router(config-ipv6-acl)# 30 permit ipv6 10:1:1::4/64 any capture
```

2. Apply the traffic monitoring to an interface.

```
Router(config)# interface HundredGigE0/1/0/3
Router(config-if)# ipv6 address 10:1:1::5/64
Router(config-if)# monitor-session mon2 ethernet direction rx-only port-level
Router(config-if-mon)# acl
Router(config-if-mon)# ipv6 access-group acl2 ingress
```

Verification

```
RP/0/RP0/CPU0:ios#show running-config ipv6 access-list acl2
Thu Aug 14 20:22:54.388 UTC
ipv6 access-list acl2
 10 permit ipv6 10:1:1::2/64 any capture
 20 permit ipv6 10:1:1::3/64 any
 30 permit ipv6 10:1:1::4/64 any capture
!
```

Local SPAN Rate Limit

Local SPAN rate limiting takes place at the session level and not at source interface level. For rate limiting, local SPAN session should configure a traffic class. This traffic class is used to shape traffic on an egress interface. A QoS policy is applied to the egress interface over which mirrored traffic is sent.

Example for Local SPAN Rate Limit Configuration

```
Router# monitor-session mon2 ethernet
destination interface HundredGigE0/1/0/19
traffic-class 5

class-map match-any TC5
match traffic-class 5
end-class-map
```

```

policy-map shape-foo
  class TC5 /* This has to match the class that was configured on monitor session */
  shape average percent 15
  class class-default

interface HundredGigE0/1/0/19 /* This is the egress interface over which mirrored packets
are sent */
  service-policy output shape-foo

```

Traffic Mirroring with DSCP

Differentiated Service Code Point (DSCP) value of Differentiated Services (DS) field in IP packet is used to classify the traffic in the network. DS field formerly known as Type of Service (ToS). You can set the DSCP value in the six most significant bits of the differentiated services (DS) field of the IP header, thereby giving $2^6 = 64$ different values (0 to 63). These six bits affect the Per Hop Behavior (PHB) and hence affects how a packet is moved forward. The default value of DSCP is zero (0). DSCP was defined under RFC 2474.

Following the principle of traffic classification, DSCP places a particular packet into a limited number of traffic classes. Similarly, the router is also informed about the DSCP values and the router can prioritize the packet in traffic flow.

Refer the table to know more about the service class names defined in [RFC 2474](#).

Table 27: DSCP, DS, and ToS values

DSCP Value in Decimal	DS Binary	DS Hex	DSCP Name	DS/ToS Value	Service Class
0	000000	0x00	DF/CS0	0	Standard
-	-	-	none	2	
1	000001	0x01	None	4	
1	000001	0x01	LE	4	Lower-effort
2	000010	0x02	None	8	
4	000100	0x04	None	16	
8	001 000	0x08	CS1	32	Low-priority data
10	001 010	0x0a	AF11	40	High-throughput data
12	001 100	0x0c	AF12	48	High-throughput data
14	001 110	0x0e	AF13	56	High-throughput data
16	010 000	0x10	CS2	64	OAM
18	010 010	0x12	AF21	72	Low-latency data
20	010 100	0x14	AF22	80	Low-latency data
22	010 010	0x16	AF23	88	Low-latency data
24	011 000	0x18	CS3	96	Broadcast video
26	011 000	0x1a	AF31	104	Multimedia streaming

28	011 100	0x1c	AF32	112	Multimedia streaming
30	011 110	0x1e	AF33	120	Multimedia streaming
32	100 000	0x20	CS4	128	Real-timeinteractive
34	100 010	0x22	AF41	136	Multimedia conferencing
36	100 100	0x24	AF42	144	Multimedia conferencing
38	100 110	0x26	AF43	152	Multimedia conferencing
40	101 000	0x28	CS5	160	Signaling(IP telephony, etc)
44	101 100	0x2c	Voice-admit	176	
46	101 110	0x2e	EF	184	Telephony
48	110 000	0x30	CS6	192	Networkrouting control
56	111 000	0x38	CS7	224	“reserved”

DSCP Marking on Egress GRE Tunnel in ERSPAN

Table 28: Feature History Table

Feature Name	Release Information	Feature Description
DSCP Marking on Egress GRE Tunnel in ERSPAN	Release 7.5.4	You can now set or modify Differentiated Service Code Point (DSCP) value on the ERSPAN GRE tunnel header. This feature allows you to control the QoS for your network's ERSPAN GRE tunnel traffic and eases the effort to control your customers' bandwidth across next-hop routers.

Starting Cisco IOS XR Software Release 7.5.4, you can set or modify the DSCP marking on the ERSPAN GRE tunnels. ERSPAN uses GRE encapsulation to route captured traffic.

Configure DSCP Marking on Egress GRE Tunnel in ERSPAN

Configuration Example

This example shows how you can configure DSCP Marking on Egress GRE tunnel in ERSPAN.

```
Router#configure terminal
Router(config)#interface tunnel-ip1
Router(config-if)#tunnel tos 96
Router(config-if)#tunnel mode gre ipv4
Router(config-if)#tunnel source 192.0.2.1
Router(config-if)#tunnel destination 192.0.2.254
Router(config-if)#commit
```



Note You can configure DSCP value on both IPv4 and IPv6 headers.

Running Configuration

```
interface tunnel-ip1
  tunnel tos 96
  tunnel mode gre ipv4
  tunnel source 192.0.2.1
  tunnel destination 192.0.2.254
!
```

Verification

You can use the following commands to verify that ToS value is configured:

```
Router#show run interface tunnel-ip 1
interface tunnel-ip1
  ipv4 address 192.0.2.0/24
  tunnel tos 96
  tunnel mode gre ipv4
  tunnel source 192.0.2.1
  tunnel vrf red
  tunnel destination 192.0.2.254

Router#show monitor-session ERSPAN-2 status internal
Information from SPAN Manager and MA on all nodes:
Monitor-session ERSPAN-2 (ID 0x00000003) (Ethernet)
SPAN Mgr: Destination interface tunnel-ip1 (0x20008024)
  Last error: Success
  Tunnel data:
    Mode: GREoIPv4
    Source IP: 192.0.2.1
    Dest IP: 192.0.2.254
    VRF: red
    VRF TBL ID: 0
    ToS: 96
    TTL: 255
    DFbit: Not set
```

DSCP Bitmask to Filter Ingress ERSPAN Traffic

Table 29: Feature History Table

Feature Name	Release Information	Feature Description
DSCP Bitmask to Filter Ingress ERSPAN Traffic	Release 7.5.4	<p>You can now mirror multiple traffic flows for matched Differentiated Service Code Point (DSCP) value of IP header on the Encapsulated remote SPAN (ERSPAN). The matched DSCP value is based on the DSCP value and the bitmask configured in Access Control List (ACL) rule.</p> <p>Earlier, you could monitor single traffic flow by setting the RFC 4594 defined DSCP values in the GRE tunnel header.</p> <p>This feature introduces the following changes:</p> <ul style="list-style-type: none"> • CLI: <code>deny (IPv4)</code>, <code>deny (IPv6)</code>, <code>permit (IPv4)</code>, and <code>permit (IPv6)</code> are modified to include new keyword bitmask. • YANG DATA Model: New XPathS for Cisco-IOS-XR-um-ipv4-access-list-cfg and Cisco-IOS-XR-um-ipv6-access-list-cfg (see Github, YANG Data Models Navigator).

Starting Release 7.5.4, You can configure an ACL rule with DSCP bitmask on the ERSPAN GRE tunnels to mirror specific traffic flows.

Without ACL rule, ERSPAN mirrors all the traffic on the incoming port. When ACL is configured with DSCP and DSCP mask on the ERSPAN, ERSPAN mirrors the traffic whose DSCP value lies within the combination of DSCP value and the specified mask.

A DSCP value is mapped to a single traffic class as per the defined value in [RFC2474](#). Masking the DSCP value in ACL rule allows to mirror multiple traffic flows. DSCP value and mask operate similar to IPv4 address and mask.

Configure DSCP Bitmask to Filter Ingress ERSPAN Traffic

To configure DSCP bitmask, use the `bitmask` option along with the `dscp` option while configuring the ACL.

Configuration Example for IPv4

This example shows how you can configure DSCP bitmask on ingress ERSPAN for IPv4 traffic.

```
/*configure the ACL*/
Router# config
Router(config)# ipv4 access-list acl1
Router(config-ipv4-acl)# 10 permit ipv4 host 192.0.2.1 any dscp af22 bitmask 0x3f
Router(config-ipv4-acl)# commit
Router(config-ipv4-acl)# exit

/* Perform the following configurations to attach the created ACL to an interface*/
Router(config)# interface HundredGigE0/0/0/6
Router(config-if)# ipv4 address 192.0.2.51 255.255.255.0
```

```
/* Monitor the ingress ACL applied and DSCP masked IPv4 traffic on ERSPAN*/
Router(config-if)# monitor-session TEST ethernet direction rx-only port-level acl ipv4 acl1
Router(config-if)# commit
```

Running Configuration

```
Router(config)# show running-config ipv4 access-list
ipv4 access-list acl1
 10 permit ipv4 host 192.0.2.1 any dscp af22 bitmask 0x3f
!

interface HundredGigE0/0/0/6
 ipv4 address 192.0.2.51 255.255.255.0
 monitor-session TEST ethernet direction rx-only port-level acl ipv4 acl1
!
!
```

Configuration Example for IPv6

This example shows how you can configure DSCP bitmask on ingress ERSPAN for IPv6 traffic.

```
/*configure the ACL*/
Router# config
Router(config)# ipv6 access-list acl1
Router(config-ipv6-acl)# 10 permit ipv6 host 2001:DB8::2/32 any dscp 33 bitmask 0x3f
Router(config-ipv6-acl)# commit
Router(config-ipv6-acl)# exit

/* Perform the following configurations to attach the created ACL to an interface*/
Router(config)# interface HundredGigE 0/0/10/3
Router(config-if)# ipv6 address 2001:DB8::1/32

/* Monitor the ingress ACL applied and DSCP masked IPv4 traffic on ERSPAN*/
Router(config-if)# monitor-session TEST ethernet direction rx-only port-level acl ipv6 acl1
Router(config-if)# commit
```

Running Configuration

```
Router(config)# show running-config ipv6 access-list
ipv6 access-list acl1
 10 permit ipv6 acl1 host 2001:DB8::2/32 any dscp 33 bitmask 0x3f
!

interface HundredGigE0/0/10/3
 ipv6 address 2001:db8::1/32
 monitor-session TEST ethernet direction rx-only port-level acl ipv6 acl1
!
!
```

Monitor Multiple ERSPAN Sessions with SPAN and Security ACL

Table 30: Feature History Table

Feature Name	Release Information	Feature Description
Monitor Multiple ERSPAN Sessions with SPAN and Security ACL	Release 7.5.4	With this feature, you can use SPAN and security ACL together to monitor multiple ERSPAN sessions under the same source interface. SPAN ACL helps you to distribute the mirrored traffic over different destination interfaces and Security ACL helps you to allow selective incoming traffic.

Starting Cisco IOS XR Software Release 7.5.4 you can monitor multiple ERSPAN sessions using GREv4 and GREv6 under the same source interface. Multiple ERSPAN monitor sessions configured on an interface allow you to choose the destination interface for the mirrored traffic. For the configuration of monitor sessions, you can use SPAN and security ACLs together. The SPAN and security ACLs are applicable only in the ingress traffic.

Configure Multiple Monitor ERSPAN Sessions with SPAN and Security ACL

This example shows how to configure SPAN and Security ACL for SPAN with GREv4 and GREv6 Monitor Sessions.

Configuration example

Use the following configuration to attach SPAN and security ACLs for traffic mirroring.

```
Router# config
/*Perform the following configurations to attach the SPAN ACL to an interface*/
Router(config-if)#monitor-session always-on-v4 ethernet direction rx-only port-level
Router(config-if-mon)#acl ipv4 v4-monitor-acl1
Router(config-if-mon)#acl ipv6 v6-monitor-acl1
Router(config-if-mon)#exit
Router(config-if)#monitor-session on-demand-v4 ethernet direction rx-only port-level
Router(config-if-mon)#acl ipv4 v4-monitor-acl2
Router(config-if-mon)#acl ipv6 v6-monitor-acl2
Router(config-if-mon)#exit
/*Perform the following configurations to attach the security ACL to an interface*/
Router(config-if)#ipv4 access-group sec_aclv4 ingress
Router(config-if)#ipv6 access-group sec_aclv6 ingress
Router(config-if)#commit
```

Running configuration

```
Router(config)#show running-config interface
monitor-session always-on-v4 ethernet direction rx-only port-level
```

```

acl ipv4 v4-monitor-acl2
acl ipv6 v6-monitor-acl2
!
monitor-session on-demand-v4 ethernet direction rx-only port-level
acl ipv4 v4-monitor-acl2
acl ipv6 v6-monitor-acl2
!
ipv4 access-group sec_aclv4 ingress
ipv6 access-group sec_aclv6 ingress
!
!
```

SPAN to File

Table 31: Feature History Table

Feature Name	Release Information	Feature Description
SPAN to File Support for TX and RX	Release 7.5.3	<p>With this feature, the ability to capture the packet in TX direction along with the ability to store the capture on the file is supported.</p> <p>You can now capture the packet in the TX direction and store the capture on the file. Earlier, you could only capture or mirror the traffic in the RX direction. You now have the flexibility to choose TX, RX, or both directions.</p> <p>You can now capture and analyze the outgoing (TX) packets.</p>
Partial Packet Capture Ability for SPAN-to-File (RX)	Release 7.5.3	<p>With this feature, you can perform partial packet capture in the RX direction.</p> <p>Earlier, the ability for entire packet capture was available in the TX direction only, now you can choose entire or partial packet capture in the Rx direction also.</p> <p>Here, partial packet capture is also known as truncation.</p>
SPAN to File - PCAPng File Format	Release 7.3.1	<p>PCAPng is the next generation of packet capture format that contains a dump of data packets captured over a network and stored in a standard format.</p> <p>The PCAPng file contains different types of information blocks, such as the section header, interface description, enhanced packet, simple packet, name resolution, and interface statistics. These blocks can be used to rebuild the captured packets into recognizable data.</p> <p>The PCAPng file format:</p> <ul style="list-style-type: none"> • Provides the capability to enhance and extend the existing capabilities of data storage over time • Allows you to merge or append data to an existing file. • Enables to read data independently from network, hardware, and operating system of the machine that made the capture.

SPAN to File is an extension of the pre-existing SPAN feature that allows network packets to be mirrored to a file instead of an interface. This helps in the analysis of the packets at a later stage. The file format is PCAP, which helps that data to be used by tools, such as tcpdump or Wireshark.



Note A maximum of 100 source ports are supported across the system. Individual platforms may support lower numbers. All the SPAN sessions are configured under the Ethernet class. At any given time, the system supports four SPAN to File sessions.

When you configure a file as a destination for a SPAN session, the system creates buffer on each node to which the network packets are logged. The buffer is for all packets on the node regardless of which interface they are from. That is, multiple interfaces can provide packets to the same buffer. The system deletes the buffer when the session configuration is removed. Each node writes a file on the active RP, which contains the node ID of the node on which the buffer was located.

The minimum buffer size is 1KB. The maximum buffer size is 1000KB and default buffer size is 2KB.

If multiple interfaces are attached to a session, then interfaces on the same node are expected to have their packets sent to the same file. Bundle interfaces can be attached to a session with a file destination, which is similar to attaching individual interfaces.

From Cisco IOS XR Software Release 7.5.3 onwards, the capture of all the outgoing packets from the router is supported.

Earlier to Cisco IOS XR Software Release 7.5.3, there was no functionality which enables to capture the payload of packets coming from your customers for security reasons.

Limitations and Restrictions for SPAN to File

- Only incoming packet mirroring on the source interface is supported. Outgoing mirrored packets cannot be dumped to the file.

However, from Cisco IOS XR Software Release 7.5.3 onwards, there are no restrictions.

- SPAN ACLs can only be applied in ingress direction only. Hence, ACLs for SPAN to File can only be applied in ingress direction only.
- ACL on MPLS traffic is not supported.
- MPLS over GRE traffic is supported, however, GRE interfaces cannot be configured as source interfaces.
- Packet truncation applies for SPAN to File and ERSPAN interfaces only. If you change the destination to Local SPAN, then an **ios_msg** is displayed as a warning. The entire packet is mirrored after this message is displayed.

Example: The Partial Packet Capture feature is not supported by Local SPAN. The entire Packet will be mirrored.

- Packet truncation is per monitor session.
- Currently, truncation per interface is not supported.
- For outgoing (TX) SPAN to File, Security ACL is not supported.
- For outgoing (TX) SPAN to File, only transit traffic is mirrored.
Self-originating traffic cannot be mirrored.

Supported capabilities for SPAN to File

- The ability to mirror outgoing traffic and punt it to the CPU across all NPU versions.
- Ability to mirror outgoing IPv4, IPv6, and MPLS traffic to file.
- Ability to mirror outgoing traffic across all types of L3 interfaces, including physical, sub, bundle, and bundle sub interfaces
- Ability to mirror outgoing traffic across L2 or BVI interfaces.
- Ability to enable the new SPAN to File truncation configuration for both RX and TX direction. You can specify the `both` keyword to enable RX and TX mirroring on a single source interface.

See [Configuring SPAN to File for Truncation and Direction, on page 185](#)

- Ability to configure a different truncation size on each monitor session.
- Ability to configure SPAN to File mirroring packet truncation size from 1 to 10000. If you try to configure a value out of the range, the configuration will not accept it and displays an error message.
- Ability to change the truncation size, when packet collecting has stopped. Removing or re-adding the monitor session is not required.
- Ability to change the truncation size during packet collecting ON. Not required to stop the monitor session.
- The entire packet is mirrored by default, without the mirror first (truncation size) configuration.

Also, if the packet size is less than the configured truncation size, the entire packet is mirrored.

Action Commands for SPAN to File

Action commands allows you to start and stop network packet collection. You can run the action commands on sessions where the destination is a file. The action command autocompletes names of the globally configured SPAN to File sessions. The following table provides more information on action commands.

Table 32: Action Commands for SPAN to File

Action	Command	Description
Start	<code>monitor-session <name></code> <code>packet-collection start</code>	Use this command to start writing packets for the specified session to the configured buffer.

Action	Command	Description
Stop	<pre>monitor-session <name> packet-collection stop [discard-data write directory <dir> filename <filename>]</pre>	<p>Use this command to stop writing packets to the configured buffer. If you specify the <code>discard-data</code> option, the system clears the buffer. If you specify the <code>write</code> option, the system writes the buffer to disk before clearing it.</p> <p>When writing the buffer to disk, save the file in <code>.pcap</code> format at the following location:</p> <pre>/<directory>/<node_id>/<filename>.</pre> <p>If you include a <code>.pcap</code> extension when specifying the filename, the system will remove it to prevent the extension from being added twice.</p>

Configuring SPAN to File

Use the following command to configure SPAN to File:

```
monitor-session <name> [ethernet|ipv4|ipv6|mpls-ipv4|mpls-ipv6]
  destination file [size <kbytes>] [buffer-type linear]
```

The `monitor-session <name> [ethernet|ipv4|ipv6|mpls-ipv4|mpls-ipv6]` part of the command creates a monitor-session with the specified name and class and is a pre-existing chain point from the current SPAN feature. The `destination file [size <kbytes>] [buffer-type linear]` part of the command adds a new “file” option to the existing “destination”.

`destination file` has the following configuration options:

- Buffer size.
- Two types of buffer:
 - Circular: Once the buffer is full, the start is overwritten.
 - Linear: Once the buffer is full, no further packets are logged.



Note The default buffer-type is circular. Only linear buffer is explicitly configurable. Changing any of the parameters (buffer size or type) recreates the session, and clears any buffers of packets.

All configuration options which are applied to an attachment currently supported for other SPAN types should also be supported by SPAN to file. This may include:

- ACLs
- Write only first X bytes of packet.
- In Cisco IOS XR Release 7.5.3, truncation per global session is supported and not per interface.



Note These options are implemented by the platform when punting the packet.

Once a session has been created, then interfaces may be attached to it using the following configuration:

```
interface GigabitEthernet 0/0/0/0
  monitor-session <name> [ethernet|ipv4|ipv6|mpls-ipv4|mpls-ipv6]
```

The attachment configuration is unchanged by SPAN to File feature.

Configuration Examples

To configure a `mon1` monitor session, use the following commands:

```
monitor-session mon1 ethernet
  destination file size 230000
  !
```

In the above example, omitting the `buffer-type` option results in default circular buffer.

To configure a `mon2` monitor session, use the following commands:

```
monitor-session mon2 ethernet
  destination file size 1000 buffer-type linear
  !
```

To attach monitor session to a physical or bundle interface, use the following commands:

```
RP/0/RSP0/CPU0:router#show run interface Bundle-Ether 1
Fri Apr 24 12:12:59.348 EDT
interface Bundle-Ether1
monitor-session ms7 ethernet|ipv4|ipv6|mpls-ipv4|mpls-ipv6]
[direction {rx-only|tx-only|both[SW(1) ]} [port-level]
acl [<acl_name>]!
```

Running Configuration

```
!! IOS XR Configuration 7.1.1.124I
!! Last configuration change at Tue Nov 26 19:29:05 2019 by root
!
hostname OC
logging console informational
!
monitor-session mon1 ethernet
  destination file size 230000 buffer-type circular
!
monitor-session mon2 ethernet
  destination file size 1000 buffer-type linear
!
interface Bundle-Ether1
monitor-session ms7 ethernet
  direction rx-only
end
```

Verification

To verify packet collection status:

```
RP/0/RP0/CPU0:router#show monitor-session status
Monitor-session mon1
Destination File - Packet collecting
=====
Source Interface      Dir.      Status
-----
```

```

Hu0/9/0/2                Rx      Operational

Monitor-session mon2
Destination File - Packet collecting
=====
Source Interface          Dir      Status
-----
BE2.1.                    Rx      Operational

```

If packet collection is not active, the following line is displayed:

```

Monitor-session mon2
Destination File - Not collecting

```

Configuring SPAN to File for Truncation and Direction

Configuring SPAN to File for Truncation

Use the **mirror first** command in monitor session configuration mode to create a SPAN to File monitor session for mirroring the packets with truncation enabled:

```

monitor-session <name> [ethernet]
destination file [size <kbytes>] [buffer-type linear|circular]
mirror first <number>

```

Once a session has been created, then interfaces may be attached to it using the following configuration:

```

interface <>
  monitor-session session-name ethernet direction rx-only|tx-only|both | acl [acl_name]

```

Configuration Examples

To configure a `mon1` monitor session, use the following commands:

```

monitor-session mon1 ethernet
  destination file
  mirror first 128
!

```

Configuring SPAN to File for Direction

Use the following command to create a SPAN to File monitor session for mirroring the packets:

```

monitor-session mon2 ethernet
  destination file
!

```

Attach the session which has been created to the interfaces using the following configuration:

```

interface <>
  monitor-session session-name ethernet direction rx-only|tx-only|
  acl [acl_name]

```

Running Configuration for all

```

monitor-session mon3 ethernet
  destination file
!

interface Hu0/9/0/2
  monitor-session mon1 ethernet
  direction rx-only
!

```

```

interface bundle-ether1
monitor-session mon2 ethernet
direction tx-only
!

interface bundle-ether2.1
monitor-session mon3 ethernet
direction both
end

```

Verification

The **show monitor-session status** displays the direction.

```

Router#show monitor-session status
Monitor-session mon1
Destination File - Packet collecting
=====
Source Interface      Dir      Status
-----
Hu0/9/0/2            Rx       Operational
Monitor-session mon2
Destination File - Packet collecting
=====
Source Interface      Dir      Status
-----
BE1                   Tx       Operational
Monitor-session mon3
Destination File - Packet collecting
=====
Source Interface      Dir      Status
-----
BE2.1                 Both    Operational

```

Mirroring Forward-Drop Packets

Table 33: Feature History Table

Feature Name	Release Information	Description
Mirroring Forward-Drop Packets	Release 7.5.4	<p>Mirroring forward-drop packets feature copies or mirrors the packets that are dropped during the forwarding process at the router ingress to a configured destination. These mirrored packets can be captured and analyzed using network monitoring tools. The analysis of dropped packets helps you understand the types of traffic that are blocked, analyze potential security threats, troubleshoot, and optimize network performance.</p> <p>This feature introduces the following changes:</p> <ul style="list-style-type: none"> • CLI: forward-drop rx • YANG Data Model: New XPath for Cisco-IOS-XR-um-monitor-session-cfg.yang (see GitHub, YANG Data Models Navigator)

In a network, packets are forwarded from one device to another until they reach their destination. However, in some cases, routers may drop packets during this forwarding process. These packets are known as forward-drop packets.

The packet drop can happen for several reasons, such as congestion on the network, errors in the packet header or payload, blocking by firewall or access control lists (ACL), and so on. These forward-drop packets are typically discarded before they can reach their intended destination, and may have to be re-transmitted by the source device. This feature supports mirroring of these forward-drop packets at the ingress (Rx direction) to another destination. When a global forward-drop session is configured for the router, the forward-drop packets at the ingress are mirrored or copied to the configured destination. You can configure the mirror destination as a file (for SPAN-to-file sessions) or an IPv4 GRE tunnel ID (for ERSPAN).

Mirroring forward-drop packets to a suitable destination for analysis can help in the following:

- **Network visibility:** By mirroring and analyzing forward-drop packets, network administrators gain better visibility into the types of traffic that are blocked by the firewalls and access control lists (ACL).
- **Threat detection:** As the original dropped packet is forwarded without any change, it helps in identifying the source of potential security threats.
- **Troubleshooting:** Analyzing forward-drop packets helps in troubleshooting network issues that may be causing the packet drop. This helps in taking proactive measures to avoid escalation of the issue.

Guidelines and Restrictions for Mirroring Forward-Drop Packets

- Only one global forward-drop session can be configured on a router.
- When traffic-class is configured under monitor-session for forward-drop, the type of service (ToS) byte of the outgoing ERSPAN packet is overwritten with the configured traffic-class value.
- In-band traffic destined to router management interface cannot be captured using this functionality.
- For ERSPAN sessions that monitor forward-drop packets, a default value of 0 is used for the encapsulation traffic class, irrespective of the DSCP value assigned for the tunnel.
- ERSPAN counters are not updated for forward-drop packets.

Configuring Forward-Drop

Perform the following tasks on the router to configure a global session for mirroring forward-drop packets:

1. Configure the tunnel mode.
2. Configure the tunnel source.
3. Configure the tunnel destination.
4. Configure a traffic mirroring session.
5. Associate a destination interface with the traffic mirroring session.
6. Run **drops packet-processing rx** command to start mirroring forward-drop packets.

This example shows how to configure a global traffic mirroring session for forward-drop packets.

```
Router(config)# interface tunnel-ip 2
Router(config-if)# tunnel mode gre ipv4
```

```

Router(config-if)# tunnel source 20.20.20.20
Router(config-if)# tunnel destination 192.1.1.3
Router(config-if)# exit
Router(config)# monitor-session mon2 ethernet
Router(config)# destination interface tunnel-ip2
Router(config)# drops packet-processing rx
Router(config)# commit

```

Running Configuration

This section shows forward-drop running configuration.

```

RP/0/RSP0/CPU0:router# show running-config
interface tunnel-ip 2
tunnel mode gre ipv4
tunnel source 20.20.20.20
tunnel destination 192.1.1.3
!
monitor-session mon2 ethernet
destination interface tunnel-ip2
drops packet-processing rx
!

```

Verification

Verify the forward-drop packets are mirrored using the **show monitor-session** command.

```

Router# show monitor-session mon2 status detail
Mon Aug 15 19:14:31.975 UTC
Monitor-session mon2
  Destination interface tunnel-ip2
  All forwarding drops:
    Direction: Rx
  Source Interfaces
  -----

```

Mirroring Buffer Drop Packets

Table 34: Feature History Table

Feature Name	Release Information	Description
Mirroring Buffer Drop Packets	Release 24.2.11	<p>The SPAN to File and ERSPAN mirroring capability is enhanced to mirror dropped packets by the Traffic Management (TM) buffer when it's full and starts dropping incoming packets. This capability allows you to retain and store a mirrored copy of the dropped packets, and work effectively even during process restarts or network failovers, providing a dependable solution for traffic monitoring.</p> <p>This feature is supported only on Cisco Silicon One P100- and Q200-based routers.</p> <p>This feature introduces the following changes:</p> <ul style="list-style-type: none"> • CLI: drops • YANG Data Model: New XPath for <code>Cisco-IOS-XR-Ethernet-SPAN-cfg.yang</code> (see GitHub, YANG Data Models Navigator)

Traffic Management (TM) buffer drops can occur for various reasons, primarily due to network congestion. The TM buffer stores packets for processing; however, if it becomes full, it can no longer hold additional packets and starts dropping the packets. This condition occurs when the rate at which the packets arrive is higher than the processing rate of the buffer. With the enhanced SPAN to File and ERSPAN mirroring capability, a mirrored copy of these dropped packets is retained.

This feature covers the following packet drop scenarios:

- The router drops packets when the line rate of incoming traffic is faster than the policer configured on the ingress interface. The reason for dropping these packets is notified as `PACKET_GOT_DROPPED_DUE_TO_EXACT_METER` (displayed as `TM_EXACT_METER_DROP`).
- The router drops packets when traffic is sent at a 100% line rate on an ingress interface for regular SPAN to File configuration. The reason for dropping these packets is notified as `PACKET_GOT_DROPPED_DUE_TO_STATISTICAL_METER` (Displayed as `TM_STATISTICAL_METER_DROP`).

Benefits of Mirroring Buffer Drop Packets

This feature ensures robust traffic analysis and network performance optimization, even during restarts and failovers:

- **Data Preservation:** During process restarts, while new packet logging is paused, all previously collected data remains intact, ensuring no critical diagnostic information is lost.
- **Control Plane Stability:** The operation of this feature remains unaffected by restarts of any control plane process, allowing for uninterrupted traffic analysis.

- Failover Assurance: For interfaces not located on the Route Processor (RP), this feature ensures a seamless failover experience. Traffic mirroring and packet capture proceed without disruption, even amidst hardware changes or network reconfigurations, safeguarding continuous network analysis.

Guidelines and Restrictions for Mirroring Buffer Drop Packets

- SPAN to File Buffer drop is a global configuration, not specific to any interface.
- Mirroring of buffer drop packets is supported only on the ingress interface.
- Only the packets dropped by the reasons TM_EXACT_METER_DROP and TM_STATISTICAL_METER_DROP are mirrored.
- Maximum of one drop session with destination as Span to File and one drop session with destination as ERSPAN can be configured globally per router.
- From Release 24.2.11, one forward-drop session and one TM buffer drop session are supported for file and GRE tunnel interface destinations.
- For ERSPAN sessions that monitor buffer drop packets, a default value of 0 is used for the encapsulation traffic class, irrespective of the DSCP value assigned for the tunnel.
- ERSPAN counters are not updated for buffer drop packets.

Configure Mirroring Buffer Drop Packets

Perform the following tasks on the router to configure a global session for mirroring buffer drop packets:

For SPAN To File destination:

1. Configure a traffic mirroring session.
2. Specify the destination as file.
3. Configure the TM buffer drop session.

For ERSPAN destination:

1. Configure the tunnel mode.
2. Configure the tunnel source.
3. Configure the tunnel destination.
4. Configure a traffic mirroring session.
5. Associate a destination interface with the traffic mirroring session.
6. Configure the TM buffer drop session.

Configuration Example

Enable Feature for SPAN To File Destination

The following example shows how to configure TM buffer drop mirroring for packets for SPAN To File Destination.


```

Router(config)# monitor-session S2F_sessnethernet /* Create a Span To File monitor session
*/
Router(config-mon)# destination file
Router(config-mon)# exit
Router(config)# interface HundredGigE0/0/0/0
Router(config-if)# monitor-session S2F_sessnethernet direction rx-only /* Attach Span To
File monitor session to the interface */
Router(config-if-mon)# exit
Router(config-if)# exit
Router(config)# monitor-session mon1 ethernet /* Create global monitor session for TM drop
packets */
Router(config-mon)# destination file
Router(config-mon)# drops traffic-management rx /* Enable TM buffer drop feature for SPAN
To File Destination */
Router(config-mon)# commit

```

Enable Feature for ERSPAN Destination

You can use the following example configuration to enable TM buffer drop mirroring for ERSPAN destination.

```

Router(config)# interface tunnel-ip2
Router(config-if)# tunnel mode gre ipv4
Router(config-if)# tunnel source 10.10.10.10
Router(config-if)# tunnel destination 192.0.2.1
Router(config-if) exit
Router(config)# monitor-session mon2 ethernet
Router(config-mon)# destination interface tunnel-ip2
Router(config-mon)# drops traffic-management rx
Router(config)# commit

```

Policer Configuration to check TM_EXACT_METER_DROP Packets

Configure a policer on the ingress interface and send the traffic at a line rate faster than the policer. The router drops the packets with the reason PACKET_GOT_DROPPED_DUE_TO_EXACT_METER (displayed as TM_EXACT_METER_DROP).

```

Router(config)# class-map match-any dscpl
Router(config-cmap)# match dscp ipv4 1
Router(config-cmap)# end-class-map
Router(config)# policy-map test-police-1R2C
Router(config-pmap)# class dscpl
Router(config-pmap-c)# police rate 100 mbps
Router(config-pmap-c-police)# exit
Router(config-pmap-c)# class class-default
Router(config-pmap-c)# police rate 100 mbps
Router(config-pmap-c-police)# exit
Router(config-pmap-c)# end-policy-map
Router(config)# interface HundredGigE0/0/0/0
Router(config-if)# service-policy input test-police-1R2C
Router(config-if)# commit

```

To Check TM_STATISTICAL_METER_DROP Packets

Configure a regular SPAN to File on the ingress interface and send the traffic at 100% line rate. The router drops the packets with the reason PACKET_GOT_DROPPED_DUE_TO_STATISTICAL_METER (Displayed as: TM_STATISTICAL_METER_DROP).

Running Configuration

Enable Feature for SPAN To File Destination

```

Router# show running-config
monitor-session S2F_sess ethernet
  destination file
!
interface HundredGigE0/0/0/0
  monitor-session S2F_sessnethernet direction rx-only
!
monitor-session mon1 ethernet
  destination file
  drops traffic-management rx
!

```

Enable Feature for ERSPAN Destination

```

Router# show running-config
interface tunnel-ip 2
  tunnel mode gre ipv4
  tunnel source 10.10.10.10
  tunnel destination 192.0.2.1
!
monitor-session mon2 ethernet
  destination interface tunnel-ip2
  drops traffic-management rx
!

```

Policer Configuration to check TM_EXACT_METER_DROP Packets

```

Router# show running-config
class-map match-any dscp1
  match dscp ipv4 1
  end-class-map
!
policy-map test-police-1R2C
  class dscp1
    police rate 100 mbps
  !
  !
  class class-default
    police rate 100 mbps
  !
!
end-policy-map
!
interface configure HundredGigE0/0/0/0
  service-policy input test-police-1R2C
!

```

Verification

Verify the buffer drop packets are mirrored using the **show monitor-session status** command.

```

Router# show monitor-session status
Monitor-session mon1
Destination File - Not collecting

```

```

SW Mirrored Packet Type Dir
-----
TM Drops                Rx

```

To Verify TM_STATISTICAL_METER_DROP and TM_EXACT_METER_DROP Packets

The Packets Accepted counter of the statistical meter dropped packets and exact meter dropped packets in **show controllers npu stats traps-all** should match the "SPAN drop" counter in **show spp node-counters**

within an acceptable range. This is because the statistics displayed are not updated in real-time; they refresh every 30 seconds from the hardware.

```
Router# show controllers npu stats traps-all instance all location 0/RP0/CPU0
Trap Type           NPU  Trap  Punt           Punt  Punt  Punt  Configured
Hardware   Policer Avg-Pkt  Packets           Packets           Packets
           Rate(pps) Level  Size  Accepted           ID  ID  Dest           VoQ  VLAN  TC  Rate(pps)
=====
TM_EXACT_METER_DROP                0   256  RPLC_CPU      208  1538  0   542
   523      NPU      N/A     25           18167604
TM_STATISTICAL_METER_DROP          0   257  RPLC_CPU      208  1538  0   542
   523      NPU      N/A    31143           18167604

Router# show spp node-counters location all | i SPAN
SPAN to File: 55724
SPAN drop: 31168
```

Introduction to File Mirroring

Prior to Cisco IOS XR Software Release 7.2.1 7.0.14, the router did not support file mirroring from active RP to standby RP. Administrators had to manually perform the task or use EEM scripts to sync files across active RP and standby RP. Starting with Cisco IOS XR Software Release 7.2.1 7.0.14, file mirroring feature enables the router to copy files or directories automatically from `/harddisk:/mirror` location in active RP to `/harddisk:/mirror` location in standby RP or RSP without user intervention or EEM scripts.

Two new CLIs have been introduced for the file mirroring feature:

- **mirror enable**

The `/harddisk:/mirror` directory is created by default, but file mirroring functionality is only enabled by executing the `mirror enable` command from configuration terminal. Status of the mirrored files can be viewed with `show mirror status` command.

- **mirror enable checksum**

The `mirror enable checksum` command enables MD5 checksum across active to standby RP to check integrity of the files. This command is optional.

Limitations

The following limitations apply to file mirroring:

- Supported only on Dual RP systems.
- Supports syncing only from active to standby RP. If files are copied into standby `/harddisk:/mirror` location, it won't be synced to active RP.
- A slight delay is observed in `show mirror` command output when mirror checksum configuration is enabled.
- Not supported on multichassis systems.

Configure File Mirroring

File mirroring has to be enabled explicitly on the router. It is not enabled by default.

```
RP/0/RSP0/CPU0:router#show run mirror
```

```
Thu Jun 25 10:12:17.303 UTC
mirror enable
mirror checksum
```

Following is an example of copying running configuration to `harddisk:/mirror` location:

```
RP/0/RSP0/CPU0:router#copy running-config harddisk:/mirror/run_config
Wed Jul  8 10:25:51.064 PDT
Destination file name (control-c to abort): [/mirror/run_config]?
Building configuration..
32691 lines built in 2 seconds (16345)lines/sec
[OK]
```

Verification

To verify the syncing of file copied to mirror directory, use the `show mirror` command.

```
RP/0/RSP0/CPU0:router#show mirror
Wed Jul  8 10:31:21.644 PDT
% Mirror rsync is using checksum, this show command may take several minutes if you have
many files. Use Ctrl+C to abort
MIRROR DIR: /harddisk:/mirror/
% Last sync of this dir ended at Wed Jul  8 10:31:11 2020
Location  |Mirrored |MD5 Checksum                               |Modification Time
-----|-----|-----|-----
run_config |yes      |76fc1b906bec4fe08ecda0c93f6c7815 |Wed Jul  8 10:25:56 2020
```

If checksum is disabled, `show mirror` command displays the following output:

```
RP/0/RSP0/CPU0:router#show mirror
Wed Jul  8 10:39:09.646 PDT
MIRROR DIR: /harddisk:/mirror/
% Last sync of this dir ended at Wed Jul  8 10:31:11 2020
Location  |Mirrored |Modification Time
-----|-----|-----
run_config |yes      |Wed Jul  8 10:25:56 2020
```

If there is a mismatch during the syncing process, use `show mirror mismatch` command to verify.

```
RP/0/RP0/CPU0:router# show mirror mismatch
Wed Jul  8 10:31:21.644 PDT
MIRROR DIR: /harddisk:/mirror/
% Last sync of this dir ended at Wed Jul  8 10:31:11 2020
Location  |Mismatch Reason      |Action Needed
-----|-----|-----
test.txt  |newly created item.  |send to standby
```

Traffic Mirroring Configuration Examples

This section contains examples of how to configure traffic mirroring:

Viewing Monitor Session Status: Example

This example shows sample output of the `show monitor-session` command with the `status` keyword:

```

RP/0/RP0/CPU0:router# show monitor-session status

Monitor-session cisco-rtp1
Destination interface HundredGigE0/5/0/38
=====
Source Interface   Dir   Status
-----
Gi0/5/0/4         Rx   Operational
Gi0/5/0/17        Rx   Operational

RP/0/RP0/CPU0:router# show monitor-session status detail

Monitor-session sess1
Destination interface is not configured
Source Interfaces
-----
HundredGigE0/0/0/0
  Direction: Rx
  ACL match: Enabled
  Portion: Full packet
  Status: Not operational (destination interface not known).
HundredGigE0/0/0/2
  Direction: Rx
  ACL match: Disabled
  Portion: First 100 bytes

RP/0/RP0/CPU0:router# show monitor-session status error

Monitor-session ms1
Destination interface HundredGigE0/2/0/15 is not configured
=====
Source Interface   Dir   Status
-----

Monitor-session ms2
Destination interface is not configured
=====
Source Interface   Dir   Status
-----

```

Monitor Session Statistics: Example

The monitor session statistics is provided in the form of packets and bytes. Use the following command to get the status:

```

RP/0/RP0/CPU0:router# show monitor-session <session_name> status internal
RP/0/RP0/CPU0:Router1#show monitor-session mon2 status internal
Wed Oct  9 19:39:30.402 UTC
Information from SPAN Manager and MA on all nodes:
Monitor-session mon2 (ID 0x00000001) (Ethernet)
SPAN Mgr: Destination interface tunnel-ip2 (0x0f000034)
Last error: Success
Tunnel data:
  Mode: GREoIPv4
  Source IP: 2.2.2.2
  Dest IP: 130.1.1.2
  VRF:
  ToS: 0 (copied)
  TTL: 255

```

```

    DFbit: Not set
0/1/CPU0: Destination interface tunnel-ip2 (0x0f000034)
Tunnel data:
  Mode: GREoIPv4
  Source IP: 2.2.2.2
  Dest IP: 130.1.1.2
  VRF:
  ToS: 0 (copied)
  TTL: 255
  DFbit: Not set

Information from SPAN EA on all nodes:
Monitor-session 0x00000001 (Ethernet)
0/1/CPU0: Name 'mon2', destination interface tunnel-ip2 (0x0f000034)
Platform, 0/1/CPU0:

  Monitor Session ID: 1

  Monitor Session Packets: 11
  Monitor Session Bytes: 1764

0/2/CPU0: Name 'mon2', destination interface tunnel-ip2 (0x0f000034)
Platform, 0/2/CPU0:

  Monitor Session ID: 1

  Monitor Session Packets: 0
  Monitor Session Bytes: 0

```

**Note**

- Currently, the system does not allow you to clear these counters.
- The counters are present on the line-card that contains the interface over which the mirrored packets are sent to the ERSPAN session destination.

If required, to clear the counters, delete and recreate the monitor session. Also, clear the counters by performing a Shut/No Shut of the tunnel interface, which triggers a Delete+Create action.

Layer 3 ACL-Based Traffic Mirroring: Example

This example shows how to configure Layer 3 ACL-based traffic mirroring:

```

RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# monitor-session ms1
RP/0/RP0/CPU0:router(config-mon)# destination tunnel-ip 1

RP/0/RP0/CPU0:router(config-mon)# commit

RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface HundredGigE/2/0/11
RP/0/RP0/CPU0:router(config-if)# ipv4 access-group span ingress
RP/0/RP0/CPU0:router(config-if)# monitor-session ms1 ethernet direction rx-only acl
RP/0/RP0/CPU0:router(config-if-mon)# commit

RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# ipv4 access-list span
RP/0/RP0/CPU0:router(config-ipv4-acl)# 5 permit ipv4 any any dscp 5 capture

```

```
RP/0/RP0/CPU0:router(config-ipv4-acl)# 10 permit ipv4 any any
RP/0/RP0/CPU0:router(config-ipv4-acl)# commit
```

Troubleshooting Traffic Mirroring

When you encounter any issue with traffic mirroring, begin troubleshooting by checking the output of the **show monitor-session status** command. This command displays the recorded state of all sessions and source interfaces:

```
Monitor-session sess1
<Session status>
=====
Source Interface  Dir  Status
-----
Gi0/0/0/0        Both <Source interface status>
Gi0/0/0/2        Both <Source interface status>
```

In the preceding example, the line marked as `<Session status>` can indicate one of these configuration errors:

Session Status	Explanation
Session is not configured globally	The session does not exist in global configuration. Check show run command output to ensure that a session with a correct name has been configured.
Destination interface <intf> is not configured	The interface that has been configured as the destination does not exist. For example, the destination interface may be configured to be a VLAN subinterface, but the VLAN subinterface may not have been yet created.
Destination interface <intf> (<down-state>)	The destination interface is not in Up state in the Interface Manager. You can verify the state using the show interfaces command. Check the configuration to see what might be keeping the interface from coming up (for example, a sub-interface needs to have an appropriate encapsulation configured).

The `<Source interface status>` can report these messages:

Source Interface Status	Explanation
Operational	Everything appears to be working correctly in traffic mirroring PI. Please follow up with the platform teams in the first instance, if mirroring is not operating as expected.
Not operational (Session is not configured globally)	The session does not exist in global configuration. Check the show run command output to ensure that a session with the right name has been configured.

Source Interface Status	Explanation
Not operational (destination interface not known)	The session exists, but it either does not have a destination interface specified, or the destination interface named for the session does not exist (for example, if the destination is a sub-interface that has not been created).
Not operational (source same as destination)	The session exists, but the destination and source are the same interface, so traffic mirroring does not work.
Not operational (destination not active)	The destination interface or pseudowire is not in the Up state. See the corresponding <i>Session status</i> error messages for suggested resolution.
Not operational (source state <down-state>)	The source interface is not in the Up state. You can verify the state using the show interfaces command. Check the configuration to see what might be keeping the interface from coming up (for example, a sub-interface needs to have an appropriate encapsulation configured).
Error: see detailed output for explanation	Traffic mirroring has encountered an error. Run the show monitor-session status detail command to display more information.

The **show monitor-session status detail** command displays full details of the configuration parameters, and of any errors encountered. For example:

```
RP/0/RP0/CPU: router#show monitor-session status detail
```

```
Monitor-session sess1
  Destination interface is not configured
  Source Interfaces
  -----
  HundredGigE0/0/0/0
    Direction: Both
    ACL match: Enabled
    Portion: Full packet
    Status: Not operational (destination interface not known)
  HundredGigE0/0/0/2
    Direction: Both
    ACL match: Disabled
    Portion: First 100 bytes
    Status: Not operational (destination interface not known). Error: 'Viking SPAN PD' detected
    the 'warning' condition 'PRM connection creation failure'.
Monitor-session foo
  Destination next-hop HundredGigE 0/0/0/0
  Source Interfaces
  -----
  HundredGigE 0/1/0/0.100:
    Direction: Both
    Status: Operating
  HundredGigE 0/2/0/0.200:
    Direction: Tx
    Status: Error: <blah>

Monitor session bar
```



```

No destination configured
Source Interfaces
-----
HundredGigE 0/3/0/0.100:
  Direction: Rx
  Status: Not operational(no destination)

```

Additional Debugging Commands

Here are additional trace and debug commands:

```
RP/0/RP0/CPU0:router# show monitor-session platform trace ?
```

```

all    Turn on all the trace
errors Display errors
events Display interesting events

```

```
RP/0/RP0/CPU0:router# show monitor-session trace ?
```

```
process Filter debug by process
```

```
RP/0/RP0/CPU0:router# debug monitor-session platform ?
```

```

all    Turn on all the debugs
errors VKG SPAN EA errors
event  VKG SPAN EA event
info   VKG SPAN EA info

```

```
RP/0/RP0/CPU0:router# debug monitor-session platform all
```

```
RP/0/RP0/CPU0:router# debug monitor-session platform event
```

```
RP/0/RP0/CPU0:router# debug monitor-session platform info
```

```
RP/0/RP0/CPU0:router# show monitor-session status ?
```

```

detail Display detailed output
errors  Display only attachments which have errors
internal Display internal monitor-session information
|      Output Modifiers

```

```
RP/0/RP0/CPU0:router# show monitor-session status
```

```
RP/0/RP0/CPU0:router# show monitor-session status errors
```

```
RP/0/RP0/CPU0:router# show monitor-session status internal
```

If there is no route to the destination IPv4 address, the status displayed for the monitor session looks like this:

```

RP/0/RP0/CPU0:Router1#show monitor-session mon2 status internal
Wed Oct  9 19:24:06.084 UTC
Information from SPAN Manager and MA on all nodes:
Monitor-session mon2 (ID 0x00000001) (Ethernet)
SPAN Mgr: Destination interface tunnel-ip2 (0x0f000034) (down)
Last error: Success
Tunnel data:
  Mode: GREoIPv4
  Source IP: 2.2.2.2
  Dest IP: 130.10.10.2
  VRF:
  ToS: 0 (copied)
  TTL: 255

```

```

DFbit: Not set
0/1/CPU0: Destination interface is not configured
Tunnel data:
  Mode: GREoIPv4
  Source IP: 2.2.2.2
  Dest IP: 130.10.10.2
  VRF:
  ToS: 0 (copied)
  TTL: 255
  DFbit: Not set

```

To verify if there is a route to the destination IPv4 address, use the following command:

```

RP/0/RP0/CPU0:Router1#show cef ipv4 130.10.10.2
Wed Oct  9 19:25:12.282 UTC
0.0.0.0/0, version 0, proxy default, default route handler, drop adjacency, internal 0x1001011
 0x0 (ptr 0x8e88d2b8) [1], 0x0 (0x8ea4d0a8), 0x0 (0x0)
Updated Oct  9 19:03:36.068
Prefix Len 0, traffic index 0, precedence n/a, priority 15
  via 0.0.0.0/32, 3 dependencies, weight 0, class 0 [flags 0x0]
  path-idx 0 NHID 0x0 [0x8e2db240 0x0]
  next hop 0.0.0.0/32
  drop adjacency

```

When a route is present, the command used in the previous example displays the following:

```

RP/0/RP0/CPU0:Router1#show cef ipv4 130.10.10.2
Wed Oct  9 19:26:06.141 UTC
130.1.1.0/24, version 20, internal 0x1000001 0x0 (ptr 0x8e88aa18) [1], 0x0 (0x8ea4dc68),
0x0 (0x0)
Updated Oct  9 19:26:02.139
Prefix Len 24, traffic index 0, precedence n/a, priority 3
  via 131.1.1.1/32, HundredGigE0/1/0/2, 2 dependencies, weight 0, class 0 [flags 0x0]
  path-idx 0 NHID 0x0 [0x8f8e2260 0x0]
  next hop 131.10.10.1/32
  local adjacency

```

The show monitor command displays the following:

```

show monitor-session mon2 status internal
Wed Oct  9 19:26:12.405 UTC
Information from SPAN Manager and MA on all nodes:
Monitor-session mon2 (ID 0x00000001) (Ethernet)
SPAN Mgr: Destination interface tunnel-ip2 (0x0f000034)
  Last error: Success
  Tunnel data:
    Mode: GREoIPv4
    Source IP: 2.2.2.2
    Dest IP: 130.10.10.2
    VRF:
    ToS: 0 (copied)
    TTL: 255
    DFbit: Not set
0/1/CPU0: Destination interface tunnel-ip2 (0x0f000034)
Tunnel data:
  Mode: GREoIPv4
  Source IP: 2.2.2.2
  Dest IP: 130.10.10.2
  VRF:
  ToS: 0 (copied)
  TTL: 255
  DFbit: Not set

Information from SPAN EA on all nodes:
Monitor-session 0x00000001 (Ethernet)
0/1/CPU0: Name 'mon2', destination interface tunnel-ip2 (0x0f000034)

```

```

Platform, 0/1/CPU0:

  Monitor Session ID: 1

  Monitor Session Packets: 0
  Monitor Session Bytes: 0

0/2/CPU0: Name 'mon2', destination interface tunnel-ip2 (0x0f000034)
Platform, 0/2/CPU0:

  Monitor Session ID: 1

  Monitor Session Packets: 0
  Monitor Session Bytes: 0

Missing ARP to the next hop to the destination
This condition is detected via this show command:
show monitor-session mon2 status internal

```

After resolving ARP for the next hop, which is done by invoking a ping command to the destination, the show command output displays the following:

```

RP/0/RP0/CPU0:Router1#show monitor-session mon2 status internal
Wed Oct  9 19:32:24.856 UTC
Information from SPAN Manager and MA on all nodes:
Monitor-session mon2 (ID 0x00000001) (Ethernet)
SPAN Mgr: Destination interface tunnel-ip2 (0x0f000034)
          Last error: Success
          Tunnel data:
            Mode: GREoIPv4
            Source IP: 2.2.2.2
            Dest IP: 130.10.10.2
            VRF:
            ToS: 0 (copied)
            TTL: 255
            DFbit: Not set
0/1/CPU0: Destination interface tunnel-ip2 (0x0f000034)
          Tunnel data:
            Mode: GREoIPv4
            Source IP: 2.2.2.2
            Dest IP: 130.10.10.2
            VRF:
            ToS: 0 (copied)
            TTL: 255
            DFbit: Not set

Information from SPAN EA on all nodes:
Monitor-session 0x00000001 (Ethernet)
0/1/CPU0: Name 'mon2', destination interface tunnel-ip2 (0x0f000034)
Platform, 0/1/CPU0:

  Monitor Session ID: 1
    Monitor Session Packets: 0
    Monitor Session Bytes: 0

0/2/CPU0: Name 'mon2', destination interface tunnel-ip2 (0x0f000034)
Platform, 0/2/CPU0:

  Monitor Session ID: 1
    Monitor Session Packets: 0
    Monitor Session Bytes: 0

```

Where to Go Next

When you have configured an Ethernet interface, you can configure individual VLAN subinterfaces on that Ethernet interface.

For information about modifying Ethernet management interfaces for the shelf controller (SC), route processor (RP), and distributed RP, see the *Advanced Configuration and Modification of the Management Ethernet Interface* later in this document.

For information about IPv6 see the *Implementing Access Lists and Prefix Lists* on Cisco IOS XR Software module in the *Cisco IOS XR IP Addresses and Services Configuration Guide*.



CHAPTER 10

Configuring Virtual Loopback and Null Interfaces

This module describes the configuration of loopback and null interfaces. Loopback and null interfaces are considered virtual interfaces.

A virtual interface represents a logical packet switching entity within the router. Virtual interfaces have a global scope and do not have an associated location. Virtual interfaces have instead a globally unique numerical ID after their names. Examples are Loopback 0, Loopback 1, and Loopback 99999. The ID is unique per virtual interface type to make the entire name string unique such that you can have both Loopback 0 and Null 0.

Loopback and null interfaces have their control plane presence on the active route switch processor (RP). The configuration and control plane are mirrored onto the standby RP and, in the event of a failover, the virtual interfaces move to the ex-standby, which then becomes the newly active RP.

Feature History for Configuring Loopback and Null Interfaces on Cisco IOS XR Software

Release	Modification
Release 7.0.11	This feature was introduced.

- [Prerequisites for Configuring Virtual Interfaces, on page 203](#)
- [Information About Configuring Virtual Interfaces, on page 203](#)
- [How to Configure Virtual Interfaces, on page 205](#)
- [Configuration Examples for Virtual Interfaces, on page 207](#)

Prerequisites for Configuring Virtual Interfaces

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs that you need for each command. If you suspect a user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Information About Configuring Virtual Interfaces

To configure virtual interfaces, you must understand the following concepts:

Virtual Loopback Interface Overview

A virtual loopback interface is a virtual interface with a single endpoint that is always up or active. Any packet that the system transmits over a virtual loopback interface is immediately received by the same interface. Loopback interfaces emulate a physical interface.

In Cisco IOS XR Software, virtual loopback interfaces perform these functions:

- Loopback interfaces can act as a termination address for routing protocol sessions. This allows routing protocol sessions to stay up even if the outbound interface is down.
- You can ping the loopback interface to verify that the router IP stack is working properly.

In applications where other routers or access servers attempt to reach a virtual loopback interface, you must configure a routing protocol to distribute the subnet assigned to the loopback address.

Packets routed to the loopback interface are rerouted back to the router or access server, and processed locally. IP packets routed out to the loopback interface but not destined to the loopback interface are dropped. Under these two conditions, the loopback interface can behave like a null interface.

Null Interface Overview

A null interface functions similarly to the null devices available on most operating systems. This interface is always up and can never forward or receive traffic; encapsulation always fails. The null interface provides an alternative method of filtering traffic. You can avoid the overhead that is involved with using access lists by directing undesired network traffic to the null interface.

The only interface configuration command that you can specify for the null interface is the **ipv4 unreachable** command. With the **ipv4 unreachable** command, if the software receives a nonbroadcast packet destined for itself that uses a protocol it does not recognize, it sends an Internet Control Message Protocol (ICMP) protocol unreachable message to the source. If the software receives a datagram that it cannot deliver to its ultimate destination because it knows of no route to the destination address, it replies to the originator of that datagram with an ICMP host unreachable message. By default, the system enables the **ipv4 unreachable** command. If we do not want ICMP to send protocol unreachable, then you need to configure using the **ipv4 icmp unreachable disable** command.

By default, the system creates the Null 0 interface during boot process and you cannot remove it. You can configure the **ipv4 unreachable** command for this interface, but most configuration is unnecessary because this interface just discards all the packets that the system sends.

Use the **show interfaces null0** command to display the Null 0 interface.

Virtual Management Interface Overview

Configuring an IPv4 virtual address enables you to access the router from a single virtual address with a management network without prior knowledge of which RP is active. An IPv4 virtual address persists across route switch processor (RP) failover situations. For this to happen, the virtual IPv4 address must share a common IPv4 subnet with a management Ethernet interface on both the RPs.

On a router where each RP has multiple management Ethernet interfaces, the virtual IPv4 address maps to the management Ethernet interface on the active RP that shares the same IP subnet.

Active and Standby RPs and Virtual Interface Configuration

The standby RP is available and in a state in which it can take over the work from the active RPs should that prove necessary. Conditions that necessitate the standby RP to become the active RP and assume the active RP's duties include:

- Failure detection by a watchdog
- Administrative command to take over
- Removal of the active RP from the chassis

If a second RP is not present in the chassis while the first is in operation, a second RP may be inserted and automatically becomes the standby RP. The standby RP may also be removed from the chassis with no effect on the system other than loss of RP redundancy.

After failover, the virtual interfaces all are present on the standby (now active) RP. Their state and configuration are unchanged and there has been no loss of forwarding (in the case of tunnels) over the interfaces during the failover. The routers use nonstop forwarding (NSF) over bundles and tunnels through the failover of the host RP.



Note The user need not configure anything to guarantee that the standby interface configurations are maintained. Protocol configuration such as tacacs source-interface, snmp-server trap-source, ntp source, logging source-interface do not use the virtual management IP address as their source by default. Use the **ipv4 virtual address use-as-src-addr** command to ensure that the protocol uses the virtual IPv4 address as its source address. Alternatively, you can also configure a loopback address with the designated or desired IPv4 address and set that as the source for protocols such as TACACS+ using the **tacacs source-interface** command.

How to Configure Virtual Interfaces

This section contains the following procedures:

Configuring Virtual Loopback Interfaces

This task explains how to configure a basic loopback interface.

Restrictions

- The IP address of a loopback interface must be unique across all routers on the network.
- That IP address must not be used by another interface on the router.
- The IP address must not be used by an interface on any other router on the network.

```
RP/0/RP0/CPU0:router# configure
/* Enters interface configuration mode and names the new loopback interface */
RP/0/RP0/CPU0:router#(config)# interface Loopback 3
/* Assigns an IP address and subnet mask to the virtual loopback interface */
```

```
RP/0/RP0/CPU0:router(config-if)# ipv4 address 172.18.189.38/32

RP/0/RP0/CPU0:router(config-if)# end
RP/0/RP0/CPU0:router(config-if)# commit
```

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

```
/* Display the configuration of the loopback interface */
RP/0/RP0/CPU0:router# show interfaces Loopback 3
```

Configuring Null Interfaces

This task explains how to configure a basic null interface.

```
/* Enters global configuration mode. */

RP/0/RP0/CPU0:router# configure

/* Enter the null 0 interface configuration mode. */

RP/0/RP0/CPU0:router#(config)# interface null 0

/* Save configuration changes. */

RP/0/RP0/CPU0:router(config-null0)# end

/* Verif the configuration of the null interface. */

RP/0/RP0/CPU0:router# show interfaces null 0
```

Configuring Virtual IPv4 Interfaces

This task explains how to configure an IPv4 virtual interface.

```
RP/0/RP0/CPU0:router# configure

/* Define an IPv4 virtual address for the management Ethernet interface. */
```



```
RP/0/RSP0/CPU0:router(config)# ipv4 virtual address 10.3.32.154/8
RP/0/RSP0/CPU0:router(config-null0)# end
or
RP/0/RSP0/CPU0:router(config-null0)# commit
```

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before
exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

This is an example for configuring a virtual IPv4 interface:

```
RP/0/RSP0/CPU0:router# configure
RP/0/RSP0/CPU0:router(config)# ipv4 virtual address 10.3.32.154/8
RP/0/RSP0/CPU0:router(config-null0)# commit
```

Configuration Examples for Virtual Interfaces

This section provides the following configuration examples:

Configuring a Loopback Interface: Example

The following example indicates how to configure a loopback interface:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Loopback 3
RP/0/RP0/CPU0:router(config-if)# ipv4 address 172.18.189.38/32
RP/0/RP0/CPU0:router(config-if)# end
Uncommitted changes found, commit them? [yes]: yes
RP/0/RP0/CPU0:router# show interfaces Loopback 3
```

```
Loopback3 is up, line protocol is up
Hardware is Loopback interface(s)
Internet address is 172.18.189.38/32
MTU 1514 bytes, BW Unknown
  reliability 0/255, txload Unknown, rxload Unknown
Encapsulation Loopback, loopback not set
Last clearing of "show interface" counters never
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 total input drops
```

```

0 drops for unrecognized upper-level protocol
Received 0 broadcast packets, 0 multicast packets
0 packets output, 0 bytes, 0 total output drops
Output 0 broadcast packets, 0 multicast packets

```

Configuring a Null Interface: Example

The following example indicates how to configure a null interface:

```

RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface Null 0
RP/0/RP0/CPU0:router(config-null0)# ipv4 unreachable
RP/0/RP0/CPU0:router(config-null0)# end
Uncommitted changes found, commit them? [yes]: yes
RP/0/RP0/CPU0:router# show interfaces Null 0

```

```

Null0 is up, line protocol is up
Hardware is Null interface
Internet address is Unknown
MTU 1500 bytes, BW Unknown
  reliability 0/255, txload Unknown, rxload Unknown
Encapsulation Null, loopback not set
Last clearing of "show interface" counters never
5 minute input rate 0 bits/sec, 0 packets/sec
5 minute output rate 0 bits/sec, 0 packets/sec
  0 packets input, 0 bytes, 0 total input drops
  0 drops for unrecognized upper-level protocol
  Received 0 broadcast packets, 0 multicast packets
  0 packets output, 0 bytes, 0 total output drops
  Output 0 broadcast packets, 0 multicast packets

```

Configuring a Virtual IPv4 Interface: Example

```

RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# ipv4 virtual address 10.3.32.154/8
RP/0/RP0/CPU0:router(config-null0)# commit

```



CHAPTER 11

Configure GRE Tunnels

Tunneling provides a mechanism to transport packets of one protocol within another protocol. This chapter describes GRE tunneling protocol.

Release	Feature(s) Added
Release 7.3.1	GRE Tunnel feature was introduced.

- [GRE Tunnels, on page 209](#)
- [Unidirectional GRE Encapsulation \(GREv4\), on page 214](#)
- [Unidirectional GRE Decapsulation \(GREv4\), on page 214](#)
- [ECMP and LAG Hashing for NVGRE Flows, on page 216](#)

GRE Tunnels

Table 35: Feature History Table

Feature Name	Release Information	Description
Disabling time-to-live (TTL) decrement at GRE encapsulation	Release 7.3.2	<p>This feature allows you to disable the time-to-live (TTL) decrement of the incoming packets. The result is that encapsulation of the original incoming packet takes place without any change in the TTL value.</p> <p>This feature avoids dropping incoming packets with a TTL value equal to one after GRE encapsulation.</p> <p>Before this release, the TTL value of incoming packets was decremented by one before GRE decapsulation.</p> <p>This feature introduces the tunnel ttl disable command.</p>

Feature Name	Release Information	Description
GRE Tunnel	Release 7.3.1	<p>Generic Routing Encapsulation (GRE) provides a simple approach to transporting packets of one protocol over another protocol using encapsulation. This capability is now extended to the Cisco 8000 Series Routers.</p> <p>This feature supports:</p> <ul style="list-style-type: none"> • Unidirectional GRE encapsulation • Unidirectional GRE decapsulation <p>And introduces the following commands:</p> <ul style="list-style-type: none"> • show interface tunnel-ip <> accounting (encap) • show interface tunnel-ip <> accounting (decap)
Outer-header hashing support for MPLSoGRE and IPoGRE traffic	Release 7.3.1	<p>This feature allows load-balancing of GRE traffic in transit routers. A transit node distributes incoming GRE traffic evenly across all available ECMP links in a GRE tunnel topology. A hashing function uses GRE outer and inner header tuples such as source IP, destination IP, protocol, and router ID to determine traffic entropy. This capability is now extended to the Cisco 8000 Series Routers.</p>

Generic Routing Encapsulation (GRE) is a tunneling protocol that provides a simple generic approach to transport packets of one protocol over another protocol by means of encapsulation. GRE encapsulates a payload, that is, an inner packet that should be delivered to a destination network inside an outer IP packet. The GRE tunnel behaves as virtual point-to-point link that has two endpoints identified by the tunnel source and tunnel destination address. The tunnel endpoints send payloads through GRE tunnels by routing encapsulated packets through intervening IP networks. Other IP routers along the way do not parse the payload (the inner packet); they only parse the outer IP packet as they forward it toward the GRE tunnel endpoint. Upon reaching the tunnel endpoint, GRE encapsulation is removed and the payload is forwarded to the packet's ultimate destination.

A tunnel configured using encapsulation mode performs encapsulation of IPv4/IPv6 payload inside the GRE header. A tunnel configured using decapsulation mode performs the opposite. Here, outer GRE header is decapsulated and the inner IPv4/IPv6/MPLS payload is forwarded to the next hop router. Both encapsulation and decapsulation tunnel interfaces collect statistics periodically. The statistics can be displayed on demand using the CLI commands `show interface tunnel-ip1 accounting` and `show policy-map type pbr address-family ipv4 statistics`. For more information, see [Unidirectional GRE Encapsulation \(GREv4\)](#), on page 214 and [Unidirectional GRE Decapsulation \(GREv4\)](#), on page 214.

To perform load-balancing of GRE traffic in transit routers, a transit node distributes incoming GRE traffic evenly across all available ECMP links in a GRE tunnel topology. Furthermore, to determine traffic entropy, a hashing function uses GRE outer and inner header tuples such as source IP, destination IP, protocol, and router ID.

GRE Encapsulation and Decapsulation Over BVI

Table 36: Feature History Table

Feature Name	Release Information	Description
GRE Encapsulation and Decapsulation Over BVI	Release 7.5.4	<p>You can now transport packets using the GRE protocol over Bridge-Group Virtual Interfaces (BVI).</p> <p>This feature uses GRE to encapsulate packets between two endpoints and transmit the encapsulated packets over a BVI interface. At the destination, the GRE packet is decapsulated.</p> <p>GRE encapsulation and decapsulation over BVI allows transmitting packets securely using network layer protocols while maintaining Layer 2 connectivity between the physical interfaces.</p>

From Cisco IOS XR Release 7.5.4, GRE packets are supported over a BVI interface. This support provides GRE encapsulation and decapsulation over the BVI interfaces.

The BVI is a virtual interface within the router that acts like a normal routed interface. The BVI does not support bridging itself, but acts as a gateway for the corresponding bridge-domain to a routed interface within the router. A BVI is associated with a single bridge domain and represents the link between the bridging and the routing domains on the router.

When using GRE over BVI, the GRE header is added to the original IP packet before it is sent to the BVI. The BVI then bridges the encapsulated packet to the destination interface, which is a BVI, physical interface, or a remote network.

When the encapsulated packet reaches its destination, the receiving interface performs GRE decapsulation, which involves removing the GRE header from the original IP packet. The resulting IP packet is then forwarded to its final destination.

For information on BVI, see the *Integrated Routing and Bridging* section in the *L2VPN Configuration Guide for Cisco 8000 Series Routers*.

Supported Features on a GRE Tunnel

GRE tunnel supports the following features:

- GRE or IP-in-IP tunnels support 16 unique source addresses. These 16 unique source addresses are repeated multiple times to configure 1000 encapsulation tunnels or 64 decapsulation tunnels.
- GRE encapsulation supports the following features:
 - IPv4/IPv6 over GRE IPv4 transport
 - MPLS PoP over GRE IPv4 transport
 - ABF (Access List Based Forwarding) v4/v6 over GRE
 - VRF (Virtual Routing and Forwarding) support over GRE
- GRE decapsulation supports the following features:
 - PBR-based GRE decapsulation configuration

- CLI-based GRE decapsulation configuration
- IPv4/IPv6 over GRE decapsulation
- MPLS/SRTE over GRE decapsulation
- A GRE tunnel in decapsulation mode has only tunnel source configured, without any tunnel destination address. This decapsulated GRE tunnel behaves like a P2MP (Point-to-multipoint) tunnel, which means that an incoming GRE packet can have any source IP address and matching destination IP address to the tunnel source configured. However, once a source IP address is used for decapsulated P2MP tunnel, it cannot be re-used with other decapsulation tunnels.
- The command `tunnel ttl disable` is supported. This command controls TTL decrement of a packet being encapsulated. After configuring this command for a tunnel interface, TTL value of incoming packet is not decremented by one, and original incoming packet is encapsulated without changing the TTL. By default, `tunnel ttl disable` isn't configured. This means that the TTL of incoming packets is decremented by one before GRE encapsulation.

For example, consider an incoming packet that had the TTL value equal to one. On GRE encapsulation, the TTL value is decremented by one and becomes zero. Therefore the router will discard the packet and send an ICMP message back to the originating host. Using this feature, you can disable TTL decrement and avoid the packet discard.

Configuration Example

```
Router#configure
Router(config)#interface tunnel-ip30016
Router(config-if)#tunnel ttl disable
Router(config-if)#commit
```

Limitations for Configuring GRE Tunnels

This list describes the limitations for configuring GRE tunnels:

- GRE tunnels configured without any decapsulation or encapsulation mode support only ERPSAN feature.
- Don't create multiple GRE/IP-in-IP tunnels with the same pair of source and destination IP address or interface name. Configure all tunnels with unique source-destination pairs. In an encapsulation or decapsulation tunnel where only either source or destination is mentioned, the source-destination pair should also be unique when compared to other encapsulation or decapsulation tunnels.
- Bi-directional GRE tunnel isn't supported.
- Routing protocols over GRE tunnels aren't supported.
- Multicast over GRE isn't supported.
- GRE KA (Keep Alive) isn't supported.
- GRE parameters such as MTU (Maximum Transmission Unit) and key functionalities aren't supported.

Configure GRE Tunnels

Configuring a GRE tunnel involves creating a tunnel interface and defining the tunnel source and destination. This example shows how to configure a GRE tunnel between source and destination. The router supports only uni-directional GRE with either encapsulation or decapsulation mode.

```
Router# configure
Router(config)# interface tunnel-ip1
Router(config-if)# ipv4 address 101.0.1.2 255.255.255.0
Router(config-if)# ipv6 address 101:0:1::2/64
Router(config-if)# tunnel mode gre ipv4 [encap | decap]
Router(config-if)# tunnel source 2.2.1.1
Router(config-if)# tunnel destination 2.2.2.1/32
Router(config-if)# commit
Router(config-if)# exit
```

To configure ABFv4/v6 over GRE:

```
router static
  address-family ipv4 unicast
    201.0.1.0/24 tunnel-ip1
  address-family ipv6 unicast
    201:0:1::0/64 tunnel-ip1

ipv4 access-list abf-gre
  1 permit ipv4 any any nexthop1 ipv4 201.0.1.2
ipv6 access-list abf6-gre
  1 permit ipv6 any any nexthop1 ipv6 201:0:1::2

interface HundredGigE0/0/0/24
  ipv4 address 24.0.1.1/24
  ipv6 address 24:0:1::1/64
  ipv4 access-group abf-gre ingress
  ipv6 access-group abf6-gre ingress
!
```

To configure MPLS PoP label over GRE:

```
router static
  address-family ipv4 unicast
    201.0.1.0/24 tunnel-ip1
  address-family ipv6 unicast
    201:0:1::0/64 tunnel-ip1


mpls static
  interface HundredGigE0/0/0/24
  lsp gre
    in-label 30501 allocate
    forward path 1 resolve-nexthop 201.0.1.2 out-label pop
!
```



Note Bi-directional GRE tunnel supports only ERSPAN.

Unidirectional GRE Encapsulation (GREv4)

A tunnel configured using encapsulation mode performs encapsulation of IPv4/IPv6 payload inside the GRE header. The following figure shows GRE encapsulation. Routers in the IP cloud have no knowledge of encapsulated IP source address or destination address.



521104

Configuration

The following example shows how to configure GRE tunnel encapsulation:


```
interface tunnel-ip1
  ipv4 address 101.0.1.1/24
  ipv6 address 101:0:1::1/64
  tunnel mode gre ipv4 encap
  tunnel source [ loopback1 | <any-ipaddress> | any-interface]
  tunnel destination [ 20.0.1.1/32 | 20.0.1.0/24 | 20.0.1.0/28]

router static
  address-family ipv4 unicast
    201.0.1.0/24 tunnel-1

router static
  address-family ipv6 unicast
    201:0:1::0/64 tunnel-1
```

Unidirectional GRE Decapsulation (GREv4)

In unidirectional GRE decapsulation, the outer GRE header is decapsulated and the inner IPv4/IPv6/MPLS payload is forwarded to the next hop router. The following figure shows GRE decapsulation. In the figure, PE1 strips off outer GRE header and inner payload is forwarded as regular IPv4/IPv6/MPLS forwarding.



521105

Configuration

There are two methods to configure GRE tunnel decapsulation:

1. CLI-based tunnel decapsulation configuration

```
interface tunnel-ip1
  ipv4 address 101.0.1.1/24
  ipv6 address 101:0:1::1/64
  tunnel mode gre ipv4 decap
  tunnel source [ loopback1 | <any-ipaddress> | any-interface]
  tunnel destination [ 20.0.1.1/32 | 20.0.1.0/24 | 20.0.1.0/28]
```

2. PBR-based tunnel decapsulation configuration

```
class-map type traffic match-all test_gre1
  match protocol gre
  match destination-address ipv4 10.0.1.2 255.255.255.255
  match source-address ipv4 10.10.10.1 255.255.255.255
  end-class-map

policy-map type pbr P1-test
  class type traffic test_gre1 decapsulate gre

vrf-policy vrf default address-family ipv4 policy type pbr input P1-test
```

ECMP and LAG Hashing for NVGRE Flows

Table 37: Feature History Table

Feature Name	Release Information	Feature Description
ECMP and LAG Hashing for NVGRE Flows	Release 7.5.2	<p>This feature allows transit routers to load balance the GRE traffic, based on GRE payload.</p> <p>A transit node distributes incoming GRE traffic across ECMP and LAG paths in a GRE tunnel topology. A hashing function uses GRE payload that consists of inner Ethernet frame with destination MAC and source MAC addresses, to derive the traffic entropy.</p> <p>ECMP and LAG hashing is enabled on Cisco 8000 series routers by default.</p>

Network Virtualization using Generic Routing Encapsulation (NVGRE) endpoints are network devices that act as interfaces between physical and virtual networks. NVGRE endpoint encapsulates Ethernet data frames to and from GRE tunnel. The encapsulated GRE packet is bridged and routed to the destination. On the destination, the NVGRE endpoint decapsulates the GRE packet to recover the original Ethernet frame. NVGRE is described in RFC 7637.

NVGRE uses the following header information for encapsulation:

Header	Parameters
Outer Ethernet Header	Destination MAC address, Source MAC address
Outer IP Header	IPv4 and IPv6 addresses as delivery protocol
GRE Header	GRE protocol type 0x6558 (transparent Ethernet)
GRE Payload	Inner Ethernet frame with Destination MAC address

For load balancing the GRE traffic, the transit router uses GRE payload that consists of inner Ethernet frame with destination MAC and source MAC addresses. The transit router derives the traffic entropy information from the GRE payload.

The hashing function considers the following parameters of GRE packets, along with Router ID, for load balancing the GRE traffic:

Header	Parameters
Outer IPv4 Header	Source IP address, Destination IP address, IP (GRE)
Outer IPv6 Header	Source IP address, Destination IP address, IP (GRE)

Inner Header	Destination MAC address, Source MAC
--------------	-------------------------------------

Restrictions for ECMP and LAG Hashing for NVGRE Flows

ECMP and LAG hashing does not support:

- Outer IPv4 header with Options field.
- Outer IPv6 header with extension headers.



CHAPTER 12

Configuring 802.1Q VLAN Interfaces

This module describes the configuration and management of 802.1Q VLAN interfaces.

The IEEE 802.1Q specification establishes a standard method for tagging Ethernet frames with VLAN membership information. It defines the operation of VLAN bridges that permit the definition, operation, and administration of VLAN topologies within a bridged LAN infrastructure.

The 802.1Q standard is intended to address the problem of how to divide large networks into smaller parts so broadcast and multicast traffic does not use more bandwidth than necessary. The standard also helps provide a higher level of security between segments of internal networks.

Feature History for Configuring 802.1Q VLAN Interfaces

Release	Modification
Release 7.0.11	This feature was introduced.
Release 7.2.12	Support for Layer 2 interfaces was introduced.

- [Prerequisites for Configuring 802.1Q VLAN Interfaces, on page 219](#)
- [Information About Configuring 802.1Q VLAN Interfaces, on page 220](#)
- [How to Configure 802.1Q VLAN Interfaces, on page 221](#)
- [Configuration Examples for VLAN Interfaces, on page 226](#)

Prerequisites for Configuring 802.1Q VLAN Interfaces

You must be in a user group associated with a task group that includes the proper task IDs. The command reference guides include the task IDs required for each command. If you suspect user group assignment is preventing you from using a command, contact your AAA administrator for assistance.

Before configuring 802.1Q VLAN interfaces, ensure that you meet the following conditions:

- You must have configured a HundredGigE interface, a FourHundredGigE interface, or an Ethernet bundle interface.

Information About Configuring 802.1Q VLAN Interfaces

To configure 802.1Q VLAN interfaces, you must understand the following concepts:

802.1Q VLAN Overview

A VLAN is a group of devices on one or more LANs that you can configure so that the devices can communicate as if they were attached to the same wire. When in fact, they are located on several different LAN segments. Because VLANs are based on logical instead of physical connections, they are flexible for user and host management, bandwidth allocation, and resource optimization.

The IEEE 802.1Q protocol standard addresses the problem of dividing large networks into smaller parts so broadcast and multicast traffic does not consume more bandwidth than necessary. The standard also helps provide a higher level of security between segments of internal networks.

The 802.1Q specification establishes a standard method for inserting VLAN membership information into Ethernet frames.

Cisco IOS XR software supports VLAN subinterface configuration on 40Gigabit, HundredGig, FourHundredGig, and bundle interfaces.

802.1Q Tagged Frames

The IEEE 802.1Q tag-based VLAN uses an extra tag in the MAC header to identify the VLAN membership of a frame across bridges. This tag is used for VLAN and quality of service (QoS) priority identification. The VLANs can be created statically by manual entry or dynamically through Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP). The VLAN ID associates a frame with a specific VLAN and provides the information that switches must process the frame across the network. A tagged frame is four bytes longer than an untagged frame and contains two bytes of Tag Protocol Identifier (TPID) residing within the type and length field of the Ethernet frame and two bytes of Tag Control Information (TCI) which starts after the source address field of the Ethernet frame.

Subinterfaces

Subinterfaces are logical interfaces that you can create on a hardware interface. These software-defined interfaces allow the segregation of traffic into separate logical channels on a single hardware interface. It also allows for the better utilization of the available bandwidth on the physical interface.

You can distinguish subinterfaces from each other by adding an extension at the end of the interface name and designation. For instance, the system indicates Ethernet subinterface 23 on the physical interface designated TenGigE 0/1/0/0, by TenGigE 0/1/0/0.23.

Before the system allows a subinterface to pass traffic, it must have a valid tagging protocol encapsulation and VLAN identifier assigned. All Ethernet subinterfaces always default to the 802.1Q VLAN encapsulation. However, you must explicitly define the VLAN identifier.

Subinterface MTU

The system inherits the subinterface maximum transmission unit (MTU) from the physical interface with an additional four bytes allowed for the 802.1Q VLAN tag.

Native VLAN

The router does not support a native VLAN. However, the equivalent functionality is accomplished using an **encapsulation** command as follows:

```
encapsulation dot1q TAG-ID
```

Layer 2 VPN on VLANs

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

The configuration model for configuring VLAN attachment circuits (ACs) is similar to the model used for configuring basic VLANs, where the user first creates a VLAN subinterface, and then configures that VLAN in subinterface configuration mode. To create an AC, you need to include the **l2transport** keyword in the **interface** command string to specify that the interface is a Layer 2 interface.

VLAN ACs support three modes of L2VPN operation:

- Basic Dot1Q AC—The AC covers all frames that are received and sent with a specific VLAN tag.
- QinQ AC— Only outer tag (s-tag) of 0x88a8 and inner tag (c-tag) of 0x8100 is supported.

Keep the following in mind when configuring L2VPN on a VLAN:

- Cisco IOS XR software supports 255 ACs per LC.

Use the **show interfaces** command to display AC information.

How to Configure 802.1Q VLAN Interfaces

This section contains the following procedures:

Configuring 802.1Q VLAN Subinterfaces

This task explains how to configure 802.1Q VLAN subinterfaces. To remove these subinterfaces, see the “Removing an 802.1Q VLAN Subinterface” section.



Tip You can programmatically configure and retrieve the VLAN interfaces and subinterfaces parameters using `openconfig-vlan.yang` OpenConfig data model. To get started with using data models, see the *Programmability Configuration Guide for Cisco 8000 Series Routers*.

```
RP/0/RP0/CPU0:router# configure
```

```
/* Enter subinterface configuration mode and specifies the interface type, location, and subinterface number.
*/
```

```
RP/0/RP0/CPU0:router(config)# interface TenGigE 0/2/0/4.10
```

- Replace the *interface-path-id* argument with one of the following instances:
- Physical Ethernet interface instance, or with an Ethernet bundle instance. Naming notation is *rack/slot/module/port*, and a slash between values is required as part of the notation.
- Ethernet bundle instance. Range is from 1 through 65535.
- Replace the *subinterface* argument with the subinterface value. Range is from 0 through 4095.
- Naming notation is *interface-path-id.subinterface*, and a period between arguments is required as part of the notation.

/* Set the Layer 2 encapsulation of an interface. */

```
RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 100
```



Note • The **dot1q vlan** command is replaced by the **encapsulation dot1q** command on the Cisco 8000 Series Router. It is still available for backward-compatibility, but only for Layer 3 interfaces.

/* Assign an IP address and subnet mask to the subinterface. */

```
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 178.18.169.23/24
```

- Replace *ip-address* with the primary IPv4 address for an interface.
- Replace *mask* with the mask for the associated IP subnet. The network mask can be specified in either of two ways:
 - The network mask can be a four-part dotted decimal address. For example, 255.0.0.0 indicates that each bit equal to 1 means that the corresponding address bit belongs to the network address.
 - The network mask can be indicated as a slash (/) and number. For example, /8 indicates that the first 8 bits of the mask are ones, and the corresponding bits of the address are network address.

/* The **exit** command is not explicitly required. */

```
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# end
or
RP/0/RP0/CPU0:router(config)# commit
```

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
- Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
- Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.

- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Example

```
RP/0/RP0/CPU0:S3(config)#interface fourHundredGigE 0/5/0/1.100
RP/0/RP0/CPU0:S3(config-subif)#ipv4 address 100.100.100.100/31
RP/0/RP0/CPU0:S3(config-subif)#encapsulation dot1q 100
RP/0/RP0/CPU0:S3(config-subif)#no shutdown
RP/0/RP0/CPU0:S3(config-subif)#commit
Mon Jul 8 23:05:01.979 PDT
RP/0/RP0/CPU0:S3(config-subif)#end
RP/0/RP0/CPU0:S3#show interfaces fourHundredGigE 0/5/0/1.100 brief
Mon Jul 8 23:05:08.784 PDT
```

Intf Name	Intf State	LineP State	Encap Type	MTU (byte)	BW (Kbps)
FH0/5/0/1.100	up	up	802.1Q	1518	400000000

```
RP/0/RP0/CPU0:S3#show interfaces brief location 0/5/CPU0 | include 802.1Q
Mon Jul 8 23:07:43.929 PDT
      FH0/5/0/1.100          up          up          802.1Q  1518  400000000
RP/0/RP0/CPU0:S3#
RP/0/RP0/CPU0:S3#"
```

Configuring an Attachment Circuit on a VLAN

Use the following procedure to configure an attachment circuit on a VLAN.

SUMMARY STEPS

1. **configure**
2. **interface [HundredGigE | TenGigE | Bundle-Ether | TenGigE] interface-path] id.subinterface l2transport**
3. **encapsulation dot1q vlan-id**
4. **end** or **commit**
5. **show interfaces [HundredGigE | TenGigE] interface-path-id.subinterface**

DETAILED STEPS

Procedure

	Command or Action	Purpose
Step 1	configure Example: RP/0/RP0/CPU0:router# configure terminal	Enters global configuration mode.
Step 2	interface [HundredGigE TenGigE Bundle-Ether TenGigE] interface-path] id.subinterface l2transport Example:	Enters subinterface configuration and specifies the interface type, location, and subinterface number.

	Command or Action	Purpose
	<pre>RP/0/RP0/CPU0:router(config)# interface TenGigE 0/1/0/0.1 l2transport</pre>	<ul style="list-style-type: none"> • Replace the argument with one of the following instances: • Physical Ethernet interface instance, or with an Ethernet bundle instance. Naming notation is <i>rack/slot/module/port</i>, and a slash between values is required as part of the notation. • Ethernet bundle instance. Range is from 1 through 65535. • Replace the <i>subinterface</i> argument with the subinterface value. Range is from 0 through 4095. • Naming notation is <i>instance.subinterface</i>, and a period between arguments is required as part of the notation. • You must include the l2transport keyword in the command string; otherwise, the configuration creates a Layer 3 subinterface rather than an AC.
Step 3	<p>encapsulation dot1q <i>vlan-id</i></p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 100</pre>	Sets the Layer 2 encapsulation of an interface.
Step 4	<p>end or commit</p> <p>Example:</p> <pre>RP/0/RP0/CPU0:router(config-if-l2)# end</pre> <p>or</p> <pre>RP/0/RP0/CPU0:router(config-if-l2)# commit</pre>	<p>Saves configuration changes.</p> <ul style="list-style-type: none"> • When you issue the end command, the system prompts you to commit changes: <pre>Uncommitted changes found, commit them before exiting(yes/no/cancel)? [cancel]:</pre> <ul style="list-style-type: none"> - Entering yes saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode. - Entering no exits the configuration session and returns the router to EXEC mode without committing the configuration changes. - Entering cancel leaves the router in the current configuration session without exiting or committing the configuration changes. <ul style="list-style-type: none"> • Use the commit command to save the configuration changes to the running configuration file and remain within the configuration session.
Step 5	<p>show interfaces [HundredGigE TenGigE] <i>interface-path-id.subinterface</i></p>	(Optional) Displays statistics for interfaces on the router.

	Command or Action	Purpose
	Example: RP/0/RP0/CPU0:router# show interfaces TenGigE 0/3/0/0.1	

Removing an 802.1Q VLAN Subinterface

This task explains how to remove 802.1Q VLAN subinterfaces that have been previously configured using the Configuring 802.1Q VLAN subinterfaces section in this module.

```
RP/0/RP0/CPU0:router# configure
```

```
/* Remove the subinterface, which also automatically deletes all the configuration applied to the subinterface.
*/
```

```
RP/0/RP0/CPU0:router(config)# no interface TenGigE 0/2/0/4.10
```

- Replace the *instance* argument with one of the following instances:
 - Physical Ethernet interface instance, or with an Ethernet bundle instance. Naming notation is *rack/slot/module/port*, and a slash between values is required as part of the notation.
 - Ethernet bundle instance. Range is from 1 through 65535.
 - Replace the *subinterface* argument with the subinterface value. Range is from 0 through 4095.

Naming notation is *instance.subinterface*, and a period between arguments is required as part of the notation.



Note Repeat to remove other VLAN subinterfaces.

```
RP/0/RP0/CPU0:router(config)# end
or
RP/0/RP0/CPU0:router(config)# commit
```

- When you issue the **end** command, the system prompts you to commit changes:

```
Uncommitted changes found, commit them before exiting(yes/no/cancel)?
[cancel]:
```

- Entering **yes** saves configuration changes to the running configuration file, exits the configuration session, and returns the router to EXEC mode.
 - Entering **no** exits the configuration session and returns the router to EXEC mode without committing the configuration changes.
 - Entering **cancel** leaves the router in the current configuration session without exiting or committing the configuration changes.
- Use the **commit** command to save the configuration changes to the running configuration file and remain within the configuration session.

Configuration Examples for VLAN Interfaces

This section contains the following example:

VLAN Subinterfaces: Example

The following example shows how to create three VLAN subinterfaces at one time:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface TenGigE 0/2/0/4.1

RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 10.0.10.1/24
RP/0/RP0/CPU0:router(config-subif)# interface TenGigE0/2/0/4.2

RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 101
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 10.0.20.1/24
RP/0/RP0/CPU0:router(config-subif)# interface TenGigE0/2/0/4.3

RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 102
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 10.0.30.1/24
RP/0/RP0/CPU0:router(config-subif)# commit
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# exit

RP/0/RP0/CPU0:router# show ethernet trunk bundle-Ether 1
Trunk                               Sub types          Sub states
VLAN trunks: 1,
  1 are 802.1Q (Ether)
Sub-interfaces: 3,
  3 are up.
802.1Q VLANs: 3,
  3 have VLAN Ids,

RP/0/RP0/CPU0:router# show vlan interface

```

Interface	St Ly	MTU	Subs	L3
Te0/2/0/4.1	802.1Q	10	up	
Te0/2/0/4.2	802.1Q	20	up	
Te0/2/0/4.3	802.1Q	30	up	

```
RP/0/RP0/CPU0:router# show vlan trunks brief

```

Summary	Up	Down	Ad-Down	BE1	Up L3	1514	1000
Te0/2/0/4	802.1Q (Ether)	up					

The following example shows how to create two VLAN subinterfaces on an Ethernet bundle:

```
RP/0/RP0/CPU0:router# configure
RP/0/RP0/CPU0:router(config)# interface bundle-ether 2
RP/0/RP0/CPU0:router(config-if)# ipv4 address 192.168.2.1/24
RP/0/RP0/CPU0:router(config-if)# exit
```

```
RP/0/RP0/CPU0:router(config)# interface bundle-ether 2.1

RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 100
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 192.168.100.1/24
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# interface bundle-ether 2.2

RP/0/RP0/CPU0:router(config-subif)# encapsulation dot1q 200
RP/0/RP0/CPU0:router(config-subif)# ipv4 address 192.168.200.1/24
RP/0/RP0/CPU0:router(config-subif)# exit
RP/0/RP0/CPU0:router(config)# commit
```




CHAPTER 13

Configure IP-in-IP Tunnels

This chapter provides conceptual and configuration information for IP-in-IP tunnels.

Table 38: Feature History for Configure Tunnels

Release 7.0.11	This feature was introduced.
Release 7.0.14	Support for the following feature was introduced in Configure Tunnels: <ul style="list-style-type: none">• Extended ACL must match on the outer header for IP-in-IP Decapsulation.

Table 39: Feature History Table


Feature Name	Release Information	Feature Description
IPv4 packets with IPv6 Outer Header	Release 7.5.3	With this release, decapsulation of IPv4 and IPv6 tunnels with IPv6 outer headers are supported. This feature helps the administrators to take advantage of the benefits of IPv6, such as improved routing and security, without having to upgrade their entire network to IPv6.

Tunneling provides a mechanism to transport packets of one protocol within another protocol. IP-in-IP tunneling refers to the encapsulation and decapsulation of an IP packet as a payload in another IP packet. Cisco 8000 Series Routers support IP-in-IP decapsulation with all possible combinations of IPv4 and IPv6; that is, IPv4 over IPv4, IPv6 over IPv4, IPv4 over IPv6, and IPv6 over IPv6. For example, an IPv4 over IPv6 refers to an IPv4 packet as a payload encapsulated within an IPv6 packet and routed across an IPv6 network to reach the destination IPv4 network, where it is decapsulated.

IP-in-IP tunneling can be used to connect remote networks securely or provide virtual private network (VPN) services.

The following example provides configurations for an IPv4 or IPv6 tunnel, with the transport VRF as the default VRF for the following simplified network topology.

Figure 12: IP-in-IP Tunnel Network Topology



Configuration Example for IPv4 Tunnel

PE1 Router Configuration	PE2 Router Configuration
<pre>interface GigabitEthernet0/0/0/0 !! Link between PE1-PE2 ipv4 address 100.1.1.1/64 ! interface GigabitEthernet0/0/0/1 !! Link between CE1-PE1 ipv4 address 20.1.1.1/24 ipv6 address 20::1/64 ! interface tunnel-ip 1 ipv4 address 10.1.1.1/24 ipv6 address 10::1/64 tunnel mode ipv4 tunnel source GigabitEthernet0/0/0/0 tunnel destination 100.1.1.2 ! router static address-family ipv4 unicast 30.1.1.0/24 tunnel-ip1 address-family ipv6 unicast 30::0/64 tunnel-ip1 ! !</pre>	<pre>interface GigabitEthernet0/0/0/0 !! Link between PE1-PE2 ipv4 address 100.1.1.2/64 ! interface GigabitEthernet0/0/0/1 !! Link between PE2-CE2 ipv4 address 30.1.1.1/24 ipv6 address 30::1/64 ! interface tunnel-ip 1 ipv4 address 10.1.1.2/24 ipv6 address 10::2/64 tunnel mode ipv4 tunnel source GigabitEthernet0/0/0/0 tunnel destination 100.1.1.1 ! router static address-family ipv4 unicast 20.1.1.0/24 tunnel-ip1 address-family ipv6 unicast 20::0/64 tunnel-ip1 ! !</pre>
CE1 Router Configuration	CE2 Router Configuration
<pre>interface GigabitEthernet0/0/0/1 !! Link between CE1-PE1 ipv4 address 20.1.1.2 255.255.255.0 ipv6 address 20::2/64 ! router static address-family ipv4 unicast 30.1.1.0/24 20.1.1.1 address-family ipv6 unicast 30::0/64 20::1 ! !</pre>	<pre>interface GigabitEthernet0/0/0/1 !! Link between CE2-PE2 ipv4 address 30.1.1.2 255.255.255.0 ipv6 address 30::2/64 ! router static address-family ipv4 unicast 20.1.1.0/24 30.1.1.1 address-family ipv6 unicast 20::0/64 30::1 ! !</pre>

Configuration Example for IPv6 Tunnel

PE1 Router Configuration	PE2 Router Configuration
<pre> interface GigabitEthernet0/0/0/0 !! Link between PE1-PE2 ipv6 address 100::1/64 ! interface GigabitEthernet0/0/0/1 !! Link between CE1-PE1 vrf RED ipv4 address 20.1.1.1/24 ipv6 address 20::1/64 ! interface tunnel-ip 1 vrf RED ipv4 address 10.1.1.1/24 ipv6 address 10::1/64 tunnel mode ipv6 tunnel source GigabitEthernet0/0/0/0 tunnel destination 100::2 ! vrf RED address-family ipv6 unicast import route-target 2:1 ! export route-target 2:1 ! address-family ipv4 unicast import route-target 2:1 ! export route-target 2:1 ! router static vrf RED address-family ipv4 unicast 30.1.1.0/24 tunnel-ip1 address-family ipv6 unicast 30::0/64 tunnel-ip1 ! ! ! </pre>	<pre> interface GigabitEthernet0/0/0/0 !! Link between PE1-PE2 ipv6 address 100::2/64 ! interface GigabitEthernet0/0/0/1 !! Link between PE2-CE2 vrf RED ipv4 address 30.1.1.1/24 ipv6 address 30::1/64 ! interface tunnel-ip 1 vrf RED ipv4 address 10.1.1.2/24 ipv6 address 10::2/64 tunnel mode ipv6 tunnel source GigabitEthernet0/0/0/0 tunnel destination 100::1 ! vrf RED address-family ipv6 unicast import route-target 2:1 ! export route-target 2:1 ! address-family ipv4 unicast import route-target 2:1 ! export route-target 2:1 ! router static vrf RED address-family ipv4 unicast 20.1.1.0/24 tunnel-ip1 address-family ipv6 unicast 20::0/64 tunnel-ip1 ! ! ! </pre>
CE1 Router Configuration	CE2 Router Configuration
<pre> interface GigabitEthernet0/0/0/1 !! Link between CE1-PE1 ipv4 address 20.1.1.2 255.255.255.0 ipv6 address 20::2/64 ! router static address-family ipv4 unicast 30.1.1.0/24 20.1.1.1 address-family ipv6 unicast 30::0/64 20::1 ! ! </pre>	<pre> interface GigabitEthernet0/0/0/1 !! Link between CE2-PE2 ipv4 address 30.1.1.2 255.255.255.0 ipv6 address 30::2/64 ! router static address-family ipv4 unicast 20.1.1.0/24 30.1.1.1 address-family ipv6 unicast 20::0/64 30::1 ! ! </pre>

- [Controlling the TTL Value of Inner Payload Header, on page 232](#)

- [IP-in-IP Decapsulation, on page 232](#)
- [ECMP Hashing Support for Load Balancing, on page 241](#)

Controlling the TTL Value of Inner Payload Header

Cisco 8000 Routers allow you to control the TTL value of inner payload header of IP-in-IP tunnel packets before it gets forwarded to the next-hop router. This feature enables a router to forward custom formed IP-in-IP stacked packets even if the inner packet TTL is 1. Therefore, this feature enables you to measure the link-state and path reachability from end to end in a network.



Note After you enable or disable the decrement of the TTL value of the inner payload header of a packet, you do not need to reload the line card.

Configuration

To disable the decrement of the TTL value of inner payload header of an IP-in-IP packet, use the following steps:

1. Enter the global configuration mode.
2. Disable the decrement of TTL value of inner payload header of an IP-in-IP packet.

Configuration Example

```
/* Enter the Global Configuration mode. */
Router# configure

/* Disable the decrement of TTL value of inner payload header of an IP-in-IP packet. */
Router(config)# hw-module profile cef ttl tunnel-ip decrement disable
Router(config)# commit
```



Note Starting from Release 7.3.3, Cisco IOS XR 8000 router supports a maximum of 16 IP-in-IP decap tunnels with unique source addresses. If 15 unique tunnel sources are configured that is rounded to 95% of the tunnel hardware resource OOR threshold level. As a result, the OOR State displays *Red* in **show controllers npu resources sipidxtbl location all** command output.

Associated Commands

- [hw-module profile cef ttl tunnel-ip decrement disable](#)

IP-in-IP Decapsulation

IP-in-IP encapsulation involves the insertion of an outer IP header over the existing IP header. The source and destination address in the outer IP header point to the endpoints of the IP-in-IP tunnel. The stack of IP headers is used to direct the packet over a predetermined path to the destination, provided the network administrator knows the loopback addresses of the routers transporting the packet. This tunneling mechanism

can be used for determining availability and latency for most network architectures. It is to be noted that the entire path from source to the destination does not have to be included in the headers, but a segment of the network can be chosen for directing the packets.


In IP-in-IP encapsulation and decapsulation has two types of packets. The original IP packets that are encapsulated are called Inner packets and the IP header stack added while encapsulation are called the Outer packets.



Note The router only supports decapsulation and no encapsulation. Encapsulation is done by remote routers.


The following topology describes a use case where IP-in-IP encapsulation and decapsulation are used for different segments of the network from source to destination. The IP-in-IP tunnel consists of multiple routers that are used to decapsulate and direct the packet through the data center fabric network.

Figure 13: IP-in-IP Decapsulation Through a Data Center Network



The following illustration shows how the stacked IPv4 headers are decapsulated as they traverse through the decapsulating routers.

Figure 14: IP Header Decapsulation



Stacked IP Header in an Encapsulated Packet

The encapsulated packet has an outer IPv4 header that is stacked over the original IPv4 header, as shown in the following illustration.

Figure 15: Encapsulated Packet

[-] Frame	
[-] EthernetII	
Preamble (hex)	fb555555555555d5
Destination MAC	62:19:88:64:E2:68
Source MAC	00:10:94:00:00:02
EtherType (hex)	<auto> Internet IP
[-] IPv4 Header	
Version (int)	<auto> 4
Header length (int)	<auto> 5
ToS/DiffServ	tos (0x00)
Total length (int)	<auto> calculated
Identification (int)	0
[-] Control Flags	
Reserved (bit)	0
DF Bit (bit)	0
MF Bit (bit)	0
Fragment Offset (int)	0
Time to live (int)	255
Protocol (int)	<auto> IP
Checksum (int)	<auto> 33492
Source	192.xx.xx.xx
Destination	127.0.0.1
Header Options	
Gateway	192.0.2.10
[-] IPv4 Header	
Version (int)	<auto> 4
Header length (int)	<auto> 5
ToS/DiffServ	tos (0x00)
Total length (int)	<auto> calculated
Identification (int)	0
[-] Control Flags	
Reserved (bit)	0

385413

Configuration

You can use the following sample configuration in the routers to decapsulate the packet as it traverses the IP-in-IP tunnel:

```
Router(config)# interface loopback 0
Router(config-if)# ipv4 address 127.0.0.1/32
Router(config-if)# no shutdown
Router(config-if)# interface tunnel-ip 10
```

```
Router(config-if)# ipv4 unnumbered loopback 1
Router(config-if)# tunnel mode ipv4 decap
Router(config-if)# tunnel source loopback 0
```

- **tunnel-ip**: configures an IP-in-IP tunnel interface.
- **ipv4 unnumbered loopback address**: enables ipv4 packet processing without an explicit address, except for loopback address.
- **tunnel mode ipv4 decap**: enables IP-in-IP decapsulation.
- **tunnel source**: indicates the source address for the IP-in-IP decap tunnel with respect to the router interface.



Note You can configure the tunnel destination only if you want to decapsulate packets from a particular destination. If no tunnel destination is configured, then all the ip-in-ip ingress packets on the configured interface are decapsulated.

Running Configuration

```
Router# show running-config interface tunnel-ip 10
...
interface tunnel-ip 10
ipv4 unnumbered loopback 1
tunnel mode ipv4 decap
```

Extended ACL to Match the Outer Header for IP-in-IP Decapsulation

Starting with Cisco IOS XR Software Release 7.0.14, extended ACL has to match on the outer header for IP-in-IP Decapsulation. Extended ACL support reduces mirrored traffic throughput. This match is based only on the IPv4 protocol, and extended ACL is applied to the received outermost IP header, even if the outer header is locally terminated.

Sample configuration:

```
Router#show running-config interface bundle-Ether 50.5
Tue May 26 12:11:49.017 UTC
interface Bundle-Ether50.5
ipv4 address 101.1.5.1 255.255.255.0
encapsulation dot1q 5
ipv4 access-group ExtACL_IPinIP ingress
ipv4 access-group any_dscpegg egress
!
```

```
Router#show access-lists ipv4 ExtACL_IPinIP hardware ingress location$
Tue May 26 12:11:55.940 UTC
ipv4 access-list ExtACL_IPinIP
10 permit ipv4 192.168.0.0 0.0.255.255 any ttl gt 150
11 deny ipv4 172.16.0.0 0.0.255.255 any fragments
12 permit ipv4 any any
```

Decapsulation Using Tunnel Source Direct

Table 40: Feature History Table

Feature Name	Release Information	Feature Description
Decapsulating Using Tunnel Source Direct	Release 7.5.3	<p>Tunnel source direct allows you to decapsulate the tunnels on any L3 interface on the router.</p> <p>You can use the tunnel source direct configuration command to choose the specific IP Equal-Cost Multipath (ECMP) links for troubleshooting, when there are multiple IP links between two devices.</p>

To debug faults in various large networks, you may have to capture and analyze the network traffic at a packet level. In datacenter networks, administrators face problems with the volume of traffic and diversity of faults. To troubleshoot faults in a timely manner, DCN administrators must identify affected packets inside large volumes of traffic. They must track them across multiple network components, analyze traffic traces for fault patterns, and test or confirm potential causes.

In some networks, IP-in-IP decapsulation is currently used in network management, to verify ECMP availability and to measure the latency of each path within a datacenter.

The Network Management System (NMS) sends IP-in-IP (IPv4 or IPv6) packets with a stack (multiple) of predefined IPv4 or IPv6 headers (device IP addresses). The destination device at each hop removes the outside header, performs a lookup on the next header, and forwards the packets if a route exists.

Using the **tunnel source direct** command, you can choose the specific IP Equal-Cost Multipath (ECMP) links for troubleshooting, when there are multiple IP links between two devices.



Tip You can programmatically configure and manage the Ethernet interfaces using `openconfig-ethernet-if.yang` and `openconfig-interfaces.yang` OpenConfig data models. To get started with using data models, see the *Programmability Configuration Guide for Cisco 8000 Series Routers*.

Guidelines and Limitations

The following guidelines are applicable to this feature.

- The **tunnel source direct** command is only compatible with 'tunnel mode decap' for IP-in-IP decapsulation.
- The source-direct tunnel is always operationally `up` unless it is administratively shut down. The directly connected interfaces are identified using the **show ip route direct** command.
- All Layer 3 interfaces that are configured on the device are supported.
- Platform can accept and program only certain number of IP addresses. The number of IP addresses depends on the make of the platform linecard (LC). Each LC can have different number of Network Processor (NP) slices and interfaces.

- Only one source-direct tunnel per address-family is supported for configuration.
- Source-direct and regular decap tunnels can't co-exist for a specific address-family. Any configuration that attempts to enable both is automatically rejected, and an error message is displayed to indicate the conflict.
- Inline modification of an existing regular decap tunnel (**tunnel source interface | IP address**) to a source-direct tunnel (**tunnel source direct**), or changing a source-direct tunnel to a regular decap tunnel, is not supported. Commit-replace may fail if the same tunnel-id is used as part of the commit-replace operation. You must delete the tunnel and recreate it.

The following functionalities are not supported for the **tunnel source direct** option.

- GRE tunneling mode.
- VRF (only default VRF is supported).
- ACL and QoS on the tunnels.
- Tunnel encapsulation.
- Tunnel NetIO DLL: Decapsulation is not supported if the packet is punted to slow path.

Configure Decapsulation Using Tunnel Source Direct

Configuration

The **tunnel source direct** configures IP-in-IP tunnel decapsulation on any directly connected IP addresses. This option is now supported only when the IP-in-IP decapsulation is used to source route the packets through the network.

This example shows how to configure IP-in-IP tunnel decapsulation on directly connected IP addresses:

```
Router# configure terminal
Router(config)#interface Tunnel4
  Router(config)#tunnel mode ipv4 decap
  Router(config)#tunnel source direct
  Router(config)#no shutdown
```

This example shows how to configure IP-in-IP tunnel decapsulation on IPv6 enabled networks:

```
Router# configure terminal
Router(config)#interface Tunnel6
  Router(config)#tunnel mode ipv6 decap
  Router(config)#tunnel source direct
  Router(config)#no shutdown
```

Verifying the Configuration

The following example shows how to verify IP-in-IP tunnel decapsulation with **tunnel source direct** option:

```
Router#show running-config interface tunnel 1
interface Tunnel1
  tunnel mode ipv6ipv6 decapsulate-any
  tunnel source direct
  no shutdown

Router#show interface tunnel 1
Tunnel1 is up    Admin State: up
MTU 1460 bytes, BW 9 Kbit
```

```
Tunnel protocol/transport IPv6/DECAPANY/IPv6
Tunnel source - direct
Tx      0 packets output, 0 bytes      Rx      0 packets input, 0 bytes
```

Configure Tunnel Destination with an Object Group

Table 41: Feature History Table

Feature Name	Release Information	Description
Configure Tunnel Destination with an Object Group	Release 7.5.4	<p>You can now assign an object group as the destination for an IP-in-IP decapsulation tunnel. With this functionality, you could configure an IPv4 or IPv6 object group consisting of multiple IPv4 or IPv6 addresses as the destination for the tunnel instead of a single IPv4 or IPv6 address. Using an object group instead of a singular IP address. This helps reduce the configuration complexity in the router by replacing the multiple tunnels with one destination with a single decapsulation tunnel that supports a diverse range of destinations</p> <p>The feature introduces these changes:</p> <ul style="list-style-type: none"> • CLI: New tunnel destination command. • YANG Data Model: New object-group option supported in Cisco-IOS-XR-um-iftunnel-cfg.yang Cisco native model (see GitHub).

In IP-in-IP Decapsulation, the router accepts a packet on a tunneled interface only when the tunnel IP address matches the source IP address of the incoming packets. With this implementation, the user needs to configure separate interface tunnels for each IP address that the router needs to receive the traffic packets. This limitation often leads to configuration overload on the router.

You can eliminate the configuration overload on the router by assigning an object group as the tunnel destination for IPv4 and IPv6 traffic types. That is, the router matches the source IP address of the incoming packet against the object group available as the tunnel destination. The decapsulation tunnel accepts the incoming traffic packets when there's a match between the packet source and the object group. Otherwise, the router drops the packets.

Restrictions

The following restrictions are applicable to the tunnel destination with an object group feature:

- GRE tunnels don't support configuring object groups as the tunnel destination.
- The router supports configuring tunnel destination with an object group only when the tunnel source is tunnel source direct.
- You can configure the object group as tunnel destination only on default VRF.
- Configuring object groups as the tunnel destination isn't applicable to tunnel encapsulation.
- Subinterfaces don't support configuring object groups as the tunnel destination.
- Configuring object groups as the tunnel destination feature is mutually exclusive with ACL and QoS features.
- The tunnel destination feature supports only IPv4 and IPv6 object groups.
- The router does not support changing tunnel configuration after its creation. Configure the tunnel source direct and tunnel destination with an object group while creating the tunnel only.

Prerequisites

- Define an object group including the network elements for the tunnel destination.
- Enable the tunnel source direct feature. For more information, see decapsulation using tunnel source direct.

Configuration Example

This section provides an example for configuring the tunnel destination with an object group:

Configuration

IPv4:

```
Router# configure
/* Configure the IPv4 object group */
Router(config)# object-group network ipv4 Test_IPv4
Router(config-object-group-ipv4)# 192.0.2.0/24
Router(config-object-group-ipv4)# 198.51.100.0/24
Router(config-object-group-ipv4)# 203.0.113.0/24
Router(config-object-group-ipv4)# commit
Router(config-object-group-ipv4)# exit

/* Enters the tunnel configuration mode */
Router(config)# interface tunnel TestIPv4

/* Configures the tunnel mode */
Router(config-if)# tunnel mode ipv4 decap

/* Configures the tunnel to accept all packets with destination address matching the IP
addresses on the router */
Router(config-if)# tunnel source direct

/* Configures the tunnel to accept all packets with destination address that are in the
specified object group */
Router(config-if)# tunnel destination object-group ipv4 Test_IPv4
```

```
Router(config-if)# no shutdown
Router(config-if)# commit
Router(config-if)# exit
```

IPv6:

```
Router# configure
/* Configure the IPv6 object group */
Router(config)# object-group network ipv6 Test_IPv6
Router(config-object-group-ipv6)# 2001:DB8::/32
Router(config-object-group-ipv6)# 2001:DB8::/48
Router(config-object-group-ipv6)# commit
Router(config-object-group-ipv6)# exit

/* Enters the tunnel configuration mode */
Router(config)# interface tunnel TestIPv6

/* Configures the tunnel mode */
Router(config-if)# tunnel mode ipv6 decap

/* Configures the tunnel to accept all packets with destination address matching the IP
addresses on the router */
Router(config-if)# tunnel source direct

/* Configures the tunnel to accept all packets with destination address that are in the
specified object group */
Router(config-if)# tunnel destination object-group ipv6 Test_IPv6

Router(config-if)# no shutdown
Router(config-if)# commit
Router(config-if)# exit
```

Running Configuration

```
Router# show running config object-group
object-group network ipv4 Test_IPv4
192.0.2.0/24
198.51.100.0/24
203.0.113.0/24
!
object-group network ipv6 Test_IPv6
2001:DB8::/32
2001:DB8::/48
!

Router# show interface tunnel TestIPv4
interface TunnelTestIPv4
 tunnel mode ipv4 decap
 tunnel source direct
 tunnel destination object-group ipv4 Test_IPv4
 no shutdown
!

Router# show interface tunnel TestIPv6
interface TunnelTestIPv6
 tunnel mode ipv6 decap
 tunnel source direct
 tunnel destination object-group ipv6 Test_IPv6
 no shutdown
!
```

Verification

```
Router# show tunnel ip ea database

----- node0_0_CPU0 -----
tunnel ifhandle 0x80022cc
tunnel source 161.115.1.2
tunnel destination address group Test_IPv4
tunnel transport vrf table id 0xe0000000
tunnel mode gre ipv4, encap
tunnel bandwidth 100 kbps
tunnel platform id 0x0
tunnel flags 0x40003400
IntfStateUp
BcStateUp
Ipv4Caps
Encap
tunnel mtu 1500
tunnel tos 0
tunnel ttl 255
tunnel adjacency flags 0x1
tunnel o/p interface handle 0x0
tunnel key 0x0, entropy length 0 (mask 0xffffffff)
tunnel QT next 0x0
tunnel platform data (nil)
Platform:
Handle: (nil)
Decap ID: 0
Decap RIF: 0
Decap Recycle Encap ID: 0x00000000
Encap RIF: 0
Encap Recycle Encap ID: 0x00000000
Encap IPv4 Encap ID: 0x4001381b
Encap IPv6 Encap ID: 0x00000000
Encap MPLS Encap ID: 0x00000000
DecFEC DecRcyLIF DecStatsId EncRcyLIF
```

ECMP Hashing Support for Load Balancing

The system inherently supports the n-tuple hash algorithm. The first inner header in the n-tuple hashing includes the source port and the destination port of UDP / TCP protocol headers.

The load balancing performs these functions:

- Incoming data traffic is distributed over multiple equal-cost connections.
- Incoming data traffic is distributed over multiple equal-cost connections member links within a bundle interface.
- Layer 2 bundle and Layer 3 (network layer) load-balancing decisions are taken on IPv4, and IPv6. If it is an IPv4 or an IPv6 payload, then an n-tuple hashing is done.
- An n-tuple hash algorithm provides more granular load balancing and used for load balancing over multiple equal-cost Layer 3 (network layer) paths. The Layer 3 (network layer) path is on a physical interface or on a bundle interface.
- The n-tuple load-balance hash calculation contains:
 - Source IP address
 - Destination IP address

- IP Protocol type
- Router ID
- Source port
- Destination port
- Input interface
- Flow-label (for IPv6 only)

User-Defined Fields for ECMP Hashing

Table 42: Feature History Table

Feature Name	Release Information	Description
User-Defined Fields for ECMP Hashing	Release 24.2.11	<p>We ensure that in cases where multiple paths are used to carry packets from source to destination, each path is utilized for this purpose and no path is over-utilized or congested. This is made possible because we now provide customized ECMP hashing fields that are used for path computation.</p> <p>Previously, the router relied on fixed packet header fields for hashing, which were not user configurable. With additional user-defined bytes considered for hashing, the granularity at which the traffic can be analyzed for ECMP load balancing increases, resulting in better load balancing and path utilization.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> • cef load-balancing fields user-data • The show cef exact-route command is modified with a new user-data keyword. • The show cef ipv4 exact-route command is modified with a new user-data keyword. • The show cef ipv6 exact-route command is modified with a new user-data keyword. <p>YANG:</p> <ul style="list-style-type: none"> • New Xpath for <code>Cisco-IOS-XR-8000-fib-platform-cfg.yang</code> (see Github, YANG Data Models Navigator).

ECMP hashing is used to distribute traffic across multiple equal-cost paths. See [ECMP Hashing Support for Load Balancing, on page 241](#) for the default static hashing algorithm details.

You can now add user-defined packet header fields for ECMP path calculation for ipv4 and ipv6 flows using the **cef load-balancing fields user-data** command. Ensure you specify these user-defined fields based on the type of traffic flow that requires load balancing. You can include the following parameters:

- **Hash header:** The hash header specifies which packet header is being considered for load balancing. You can enable any or all of the available six profiles.
 - IPv4: tcp, udp, non-tcp-udp
 - IPv6: tcp, udp, non-tcp-udp

If any hash header profile is defined for load balancing, along with the fixed fields considered for hashing, additional bytes in the payload are also used for path computation.

- **Hashing offset:** The hashing offset specifies the byte location from the end of the configured header.
- **Hash size:** The hash size specifies the number of bytes that is considered from the start of the hash offset by the ECMP hashing algorithm. Range is 1 to 4 bytes.
- **Location:** This specifies the location of the ingress line card that receives the incoming traffic. The user-defined hashing configuration is applied on the specified line card.

The addition of the user-defined packet header fields increases the granularity at which the traffic is analyzed for ECMP load balancing. When multiple paths with equal cost are available for routing a specific type of packet from a source to a destination, this granularity ensures that the intended type of traffic is evenly distributed across these paths. This ensures all available paths are used efficiently and prevents congestion or over-utilization of a single path.

You can also retrieve the exact-route information based on the configured user-data using the **show cef exact-route** command with **user-data** keyword.



Note

- When the user-defined hashing configuration is active, any additional options or optional keywords are disregarded during the parsing of incoming packets for retrieving the user-defined bytes.
- The hashing results based on user-defined hash feature is applicable to BGP/IGP ECMP and LAG hashing.
- The use of the user-defined hashing configuration changes the load balancing behavior of GRE and IPinIP traffic. This includes all traffic that begins with ipv4, ipv6, ipv4+udp, ipv6+udp, ipv4+tcp, and ipv6+tcp, regardless of the payload.

Configure User-Defined Fields for ECMP Hashing

The command **cef load-balancing fields user-data** configures the additional user-defined fields that are to be considered for the hashing algorithm.

This example shows how to configure the additional IPv4 header fields for TCP packets:

```
Router# configure terminal
Router(config)# cef load-balancing fields user-data ipv4 tcp offset 5 size 3 location 0/0/CPU0
Router(config)# commit
```

- offset 5: The payload considered for hashing starts from byte 6 from the end of TCP header.
- size 3: Three bytes of payload are considered.
- location 0/0/CPU0: Specifies the line card on which the configuration is applied.

In the above example, the sixth, seventh, and eighth bytes of the payload are considered additionally for the hashing.

This example shows how to configure the additional IPv6 header fields for UDP packets:

```
Router# configure terminal
Router(config)#cef load-balancing fields user-data ipv6 udp offset 0 size 2 location 0/0/CPU0
Router(config)#commit
```

- offset 0: The payload considered for hashing starts from the end of UDP header.
- size 2: Two bytes of payload are considered.
- location 0/0/CPU0: Specifies the line card on which the configuration is applied.

In the above example, the first two bytes of payload of a UDP packet are considered additionally for the hashing.

Running Configuration

The following example shows the running configuration:

```
Router#show running-config | include cef
Fri Jul 28 12:02:01.002 UTC
cef load-balancing fields user-data ipv4 tcp offset 5 size 3 location 0/0/CPU0
cef load-balancing fields user-data ipv6 udp offset 0 size 2 location 0/0/CPU0
Router#
```

Verification

The following example shows the difference in load balancing before and after applying user-defined hashing, for a flow with data that exhibits good hashing behavior.

Before applying user-defined hashing

```
Router#show interfaces accounting | i IPV6_U
Protocol          Pkts In      Chars In      Pkts Out      Chars Out
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      2            144            0              0
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      0              0            3979416        1981749168
IPV6_UNICAST      4191438       2087336124    0              0
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      1             72             0              0
IPV6_UNICAST      1             72             0              0
Router#
```

After applying user-defined hashing

```
Router#show interfaces accounting | i IPV6_U
Protocol          Pkts In      Chars In      Pkts Out      Chars Out
IPV6_UNICAST      0              0             39119         19481262
IPV6_UNICAST      0              0             39801         19820898
IPV6_UNICAST      0              0             40483         20160534
```

IPV6_UNICAST	0	0	40524	20180952
IPV6_UNICAST	0	0	40573	20205354
IPV6_UNICAST	0	0	40614	20225772
IPV6_UNICAST	0	0	39368	19605264
IPV6_UNICAST	0	0	40734	20285532
IPV6_UNICAST	0	0	40777	20306946
IPV6_UNICAST	0	0	40171	20005158
IPV6_UNICAST	0	0	40858	20347284
IPV6_UNICAST	0	0	40269	20053962
IPV6_UNICAST	0	0	41603	20718294
IPV6_UNICAST	0	0	40363	20100774
IPV6_UNICAST	0	0	40407	20122686
IPV6_UNICAST	0	0	41098	20466804
IPV6_UNICAST	850393	423495714	0	0

To view the exact route information allocated to the packets, use **show cef exact-route** command with **user-data** keyword.

The packet contains value 0x2 in the packet position for the ipv6 packet, for which the user-defined configuration has been added for a non-tcp-udp ipv6 flow.

```
Router#show cef ipv6 exact-route 100::10 60::1 flow-label 0 protocol 59 source-port 0
destination-port 0 user-data 0x2 ingress-interface HundredGigE0/0/0/2 location 0/0/cpu0
Unsupported protocol value 59
60::/16, version 1293, internal 0x1000001 0x20 (ptr 0x8b78ef00) [1], 0x400 (0x8e9cfc48),
0x0 (0x0)
Updated Aug 14 07:50:20.022
local adjacency to Bundle-Ether3.30

Prefix Len 16, traffic index 0, precedence n/a, priority 2
  via Bundle-Ether3.30
  via fe80::72b3:17ff:feae:d703/128, Bundle-Ether3.30, 7 dependencies, weight 0, class 0
[flags 0x0]
  path-idx 7 NHID 0x0 [0x8db8bed8 0x0]
  next hop fe80::72b3:17ff:feae:d703/128
  local adjacency
```




CHAPTER 14

Configuring Generic UDP Encapsulation

Read this section to get an overview and know how to configure the Generic UDP Encapsulation.

Table 43: Feature History Table

Feature Name	Release Information	Feature Description
Outer IP Header-Driven Hash Computation for Incoming GUE Packets	Release 7.11.1	<p>We now offer you the flexibility of using only the outer IP header to calculate the hashing for incoming Generic UDP Encapsulation (GUE) packets. On enabling this feature, only the outer IP source and destination addresses are used for hashing calculations. The inner IP addresses are not considered, providing a simpler method of distribution. Previously, both inner IP and outer IP headers were used for ECMP hashing the incoming GUE packets.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> • hw-module profile gue underlay-hash <p>YANG Data Models:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-npu-hw-profile-cfg.yang</code> (see GitHub, YANG Data Models Navigator) <p>The command is supported on Q200-based ASICs.</p>

Feature Name	Release Information	Feature Description
Generic UDP Encapsulation	Release 7.3.1	<p>This feature enables you to add an additional header to packets to identify or authenticate the data using UDP. Encapsulating packets in UDP leverages the use of the UDP source port to provide entropy to Equal Cost Multi-Path (ECMP) hashing. It provides significant performance benefits for load-balancing.</p> <p>This command is introduced for this feature:</p> <p>decapsulate gue</p>

- [Understand Generic UDP Encapsulation, on page 248](#)
- [Flexible Assignment of UDP Port Numbers for Decapsulation, on page 258](#)

Understand Generic UDP Encapsulation

UDP encapsulation is a technique of adding network headers to packets and then encapsulating the packets within the User Datagram Protocol (UDP).

Encapsulating packets using UDP facilitates efficient transport across networks. By leveraging Receive Side Scaling (RSS) and Equal Cost Multipath (ECMP) routing, UDP provides significant performance benefits for load-balancing. The use of the UDP source port provides entropy to ECMP hashing and provides the ability to use the IP source or destination, and the L4 Port for load-balancing entropy.


Traditional mechanisms like Generic Routing Encapsulation (GRE) can handle only the outer Source IP address and parts of the destination address. They may not provide sufficient load balancing entropy.

Generic UDP Encapsulation (GUE) is a UDP-based network encapsulation protocol that encapsulates IPv4 and IPv6 packets. GUE provides native UDP encapsulation and defines an additional header, which helps to determine the payload carried by the IP packet. The additional header can include items, such as a virtual networking identifier, security data for validating or authenticating the GUE header, congestion control data, and so on.


In GUE, the payload is encapsulated in an IP packet that can be IPv4 or IPv6 Carrier. The UDP header is added to provide extra hashing parameters, and optional payload demultiplexing. At the decapsulation node, the Carrier IP and UDP headers are removed, and the packet is forwarded based on the inner payload.

A GUE packet has the general format:

Figure 16: GUE Packet Format



For example, in this scenario, if the data stream is sent from Host 1 to Host 2. The server acts as a GUE encapsulator that sends the packets from Host 1. The server, on the other end receiving the data, validates the data for the valid carrier IP and UDP header and decapsulates the data.



GUE has various variants, but variant 1 of GUE allows direct encapsulation of IPv4 and IPv6 in UDP. This technique saves encapsulation overhead on links for the use of IP encapsulation, and also need not allocate a separate UDP port number for IP-over-UDP encapsulation.

Variant 1 has no GUE header, but a UDP packet carries an IP packet. The first two bits of the UDP payload is the GUE variant field and match with the first two bits of the version number in the IP header.

Benefits of using GUE

- Allows direct encapsulation of payloads, such as IPv4 and IPv6 in the UDP packet.
 - You can use UDP port for demultiplexing payloads.
 - You can use a single UDP port, allowing systems to employ parsing models to identify payloads.
- Leverages the UDP header for entropy labels by encoding a tuple-based source port.

- Leverages source IP addresses for load-balance encoding. The destination too could be terminated based on a subnet providing additional bits for entropy.
- Avoids special handling for transit nodes because they only see an IP-UDP packet with some payload..
- Eases implementation of UDP tunneling with GUE. This is because of the direct encapsulation method of the payloads into UDP.

Restrictions

- Supports Generic UDP Decapsulation for only variant 1.
- Receives IPv4 packets with the defined GUE port of 6080.
- Decapsulates IPv6 packets with the defined GUE port of 6080.
- Receives MPLS packets with the UDPO MPLS port of 6635.
- Range of source or destination ports is not supported.
- Range, Source, or Destination addresses are not supported, but subnet mask entries are allowed.
- To perform decapsulation, a destination Port is mandatory.
- Terminating GRE after GUE or GUE after GRE is not supported.
- Terminating a label such as a VPN Deaggregation after GUE termination is not supported.
- Slow path support is not supported. To resolve the inner IP Adjacency, use the **cef proactive-arp-nd enable** command.
- Running the **clear all** command doesn't clear the interface of all its existing configurations.



Note

To use only outer IP header (L3 and L4) for calculating the hashing for incoming GUE packets, use the **hw-module profile gue underlay-hash enable** command. Otherwise, by default, both outer IP header (L3 and L4) and inner IP header (L3 and L4) are considered for calculating the hashing for incoming GUE packets.

The **hw-module profile gue underlay-hash enable** command is currently not supported on the P100-based and Q100-based ASICs.

Configure GUE

Table 44: Feature History Table

Feature Name	Release Information	Feature Description
Generic UDP Decapsulation for IPv6 Traffic	Release 24.1.1	<p>We've now made it possible for you to decapsulate GUEv6 packets by adding an additional header to packets that identifies or authenticates the data by using User Datagram Protocol (UDP).</p> <p>In GUE, the payload is encapsulated in an IP packet that can be an IPv6 carrier. The UDP header is added to provide extra hashing parameters, and optional payload demultiplexing. At the decapsulation node, the Carrier IP and UDP headers are removed, and the packet is forwarded based on the inner payload.</p>

Configuration Guidelines and Restrictions for GUE

- GUE for IPv6 traffic is only supported on the Cisco 8202-32FH-M router, and 88-LC0-36FH-M and 88-LC0-36FH line cards.
- GUE IPv6 decapsulation is only supported on Layer 3 ports.
- GUE IPv6 is not supported over BVI interfaces.

Configuring GUE

Use the following configuration workflow to configure GUE:

1. Configure separate GUE decap tunnel UDP destination port numbers for IPv4, IPv6, and MPLS using **hw-module profile gue udp-dest-port** command.
2. Configure a traffic class: Create a traffic class and specify various criteria for classifying packets using the match commands, and an instruction on how to evaluate these match commands.
3. Configure a policy map: Define a policy map and associate the traffic class with the traffic policy.
4. Apply the policy for each VRF, and apply this policy on all the interfaces that are part of the VRF.

Configuration Example for GUE IPv4

1. Configure separate UDP port numbers for IPv4, IPv6, and MPLS using **hw-module profile gue udp-dest-port** command.

```
Router# configure
Router# hw-module profile gue udp-dest-port ipv4 6080 ipv6 6080 mpls 6635
Router# commit
```



Note While adding or removing the **hw-module profile gue udp-dest-port** command, you must reload the router.

2. Configure a traffic class:

```

Router# configure
Router(config)# class-map type traffic match-all udp-v4
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-port 6080
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all udp-mpls1
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-port 6635
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all udp-v6
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-port 6080
Router(config-cmap)# end-class-map
Router(config)# commit

```

3. Define a policy map, and associate the traffic class with the traffic policy:

```

Router(config)# policy-map type pbr magic-decap

Router(config-pmap)# class type traffic udp-v4
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-v6
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-mpls1
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic class-default
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit
Router(config)# exit

```

4. Apply the policy for each VRF:

```

Router# configure
Router(config)# vrf-policy
Router(config-vrf-policy)# vrf default address-family ipv4 policy type pbr input magic-decap
Router(config-vrf-policy)# commit

```

Running Configuration:

```

class-map type traffic match-all udp-v4
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp
  match destination-port 6080

```

```

end-class-map
!
class-map type traffic match-all udp-v6
match destination-address ipv4 220.100.20.0 255.255.255.255
match source-address ipv4 210.100.20.0 255.255.255.255
match protocol udp
match destination-port 6080
end-class-map
!
class-map type traffic match-all udp-mpls1
match destination-address ipv4 220.100.20.0 255.255.255.255
match source-address ipv4 210.100.20.0 255.255.255.255
match protocol udp
match destination-port 6635
end-class-map
!
policy-map type pbr magic-decap
class type traffic udp-v4
  decapsulate gue variant 1
!
class type traffic udp-v6
  decapsulate gue variant 1
!
class type traffic udp-mpls1
  decapsulate gue variant 1
!
class type traffic class-default
!
end-policy-map
!

vrf-policy
  vrf default address-family ipv4 policy type pbr input magic-decap
!

```

Verification

To view the set of counter values accumulated for the packets that match the class-map:

```
Router# show policy-map type pbr addr-family ipv4 statistics
```

```

VRF Name:      default
Policy-Name:   pmap
Policy Type:   pbr
Addr Family:   IPv4

Class:      cmap-loop1
  Classification statistics      (packets/bytes)
  Matched          :              0/0
  Transmitted statistics      (packets/bytes)
  Total Transmitted  :              0/0

Class:      cmap-loop6
  Classification statistics      (packets/bytes)
  Matched          :              0/0
  Transmitted statistics      (packets/bytes)
  Total Transmitted  :              0/0

Class:      cmap-loop2
  Classification statistics      (packets/bytes)
  Matched          :              0/0
  Transmitted statistics      (packets/bytes)
  Total Transmitted  :              0/0

```

```

Class:      cmap-loop3
  Classification statistics      (packets/bytes)
    Matched                      :      198325306/17849277540
  Transmitted statistics        (packets/bytes)
    Total Transmitted            :      198325306/17849277540

Class:      cmap-loop4
  Classification statistics      (packets/bytes)
    Matched                      :                      0/0
  Transmitted statistics        (packets/bytes)
    Total Transmitted            :                      0/0

```

To clear the policy-map counters for each class-map rule, use the **clear vrf** command:

```
Router# clear vrf default address-family ipv4 statistics
```

Configuration Example for GUE IPv6

1. Configure separate UDP port numbers for IPv4, IPv6, and MPLS using **hw-module profile gue udp-dest-port** command.

```

Router# configure
Router# hw-module profile gue udp-dest-port ipv4 6080 ipv6 6080 mpls 6635
Router# commit

```



Note While adding or removing the **hw-module profile gue udp-dest-port** command, you must reload the router.

2. Configure a traffic class:

```

Router# configure
Router(config)# class-map type traffic match-all guev6_v4
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv6 11:1:1::1/128
Router(config-cmap)# match destination-port 6080
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all guev6_v6
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv6 11:1:1::6/128
Router(config-cmap)# match destination-port 6080
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all guev6_mpls
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv6 11:1:1::b/128
Router(config-cmap)# match destination-port 6635
Router(config-cmap)# end-class-map
Router(config)# commit

```

3. Define a policy map, and associate the traffic class with the traffic policy:

```

Router(config)# policy-map type pbr guev6_decap
Router(config-pmap)# class type traffic guev6_v4
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

```



```

Router(config-pmap)# class type traffic guev6_v6
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic guev6_mpls
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic class-default
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit
Router(config)# exit

```

4. Apply the policy for each VRF:

```

Router# configure
Router(config)# vrf-policy
Router(config-vrf-policy)# vrf default address-family ipv6 policy type pbr input guev6_decap
Router(config-vrf-policy)# commit

```

Running Configuration:

```

class-map type traffic match-all guev6_v4
  match protocol udp
  match destination-address ipv6 11:1:1::1/128
  match destination-port 6080
  end-class-map
!
class-map type traffic match-all guev6_v6
  match protocol udp
  match destination-address ipv6 11:1:1::6/128
  match destination-port 6080
  end-class-map
!
class-map type traffic match-all guev6_mpls
  match protocol udp
  match destination-address ipv6 11:1:1::b/128
  match destination-port 6635
  end-class-map
!

policy-map type pbr guev6_decap
  class type traffic guev6_v4
    decapsulate gue variant 1
  !
  class type traffic guev6_v6
    decapsulate gue variant 1
  !
  class type traffic guev6_mpls
    decapsulate gue variant 1
  !

vrf-policy
  vrf default address-family ipv6 policy type pbr input guev6_decap
!

```

Verification

To view the set of counter values accumulated for the packets that match the class-map:

```

Router# show policy-map type pbr vrf default addr-family ipv6 statistics

```

```

VRF Name:      default
Policy-Name:   guev6_decap
Policy Type:   pbr
Addr Family:   IPv6

Class:    guev6_v4
  Classification statistics (packets/bytes)
  Matched      :      190/24320
  Transmitted statistics (packets/bytes)
  Total Transmitted :      190/24320

Class:    guev6_v6
  Classification statistics (packets/bytes)
  Matched      :      175/24150
  Transmitted statistics (packets/bytes)
  Total Transmitted :      175/24150

Class:    guev6_mpls
  Classification statistics (packets/bytes)
  Matched      :      160/22720
  Transmitted statistics (packets/bytes)
  Total Transmitted :      160/22720

Class:    class-default
  Classification statistics (packets/bytes)
  Matched      :           0/0
  Transmitted statistics (packets/bytes)
  Total Transmitted :           0/0

```

Outer IP Header-Driven Hash Computation for Incoming GUE Packets

When multiple paths with the same cost are available for forwarding traffic, ECMP hashing is used to determine the path to select for each packet. Each packet that needs to be forwarded is processed using a hashing algorithm. The hashing algorithm considers specific packet fields such as source IP, destination IP, source port, and destination port, and generates a hash value. The generated hash value is then mapped to one of the available paths. The selected path is used to forward the packet to its destination. The goal is to distribute the traffic evenly across the available paths to prevent congestion and utilize the network resources efficiently.

Now you can use only the outer IP header (L3 and L4) for calculating the hash value for incoming GUE packets and completely ignore the usage of the inner IP header. This functionality is configurable using the CLI command **hw-module profile gue underlay-hash**. This is supported for both GUE termination (decapsulation) and GUE transit (pass-through) nodes. By default, the feature is disabled; that is, both outer IP header (L3 and L4) and inner IP header (L3 and L4) are used for calculating the hashing for GUE packets.

Benefits

- **Load Balancing Efficiency:** By hashing only on the outer IP and L4 information, the packets with the same source and destination IP addresses and L4 ports consistently follow the same path in a load-balanced environment. This helps maintain session affinity or stickiness, as the inner IP addresses or L4 port numbers may change dynamically within the encapsulated packets.
- **Network Security:** Ignoring the inner IP helps preserve privacy and confidentiality within the encapsulated packets. By focusing on the outer IP and L4 headers, the network device does not have visibility into the inner IP addressing scheme or the specific content encapsulated within the packet, which enhances security.
- **Network Scalability:** Ignoring the inner IP reduces the complexity and overhead of packet processing, improving overall network performance and scalability, especially in high-throughput environments.

Configure Outer IP Header-Driven Hash Computation for Incoming GUE Packets

This section describes how to configure hashing with only outer IP for GUE packets.

Configuration Example

Use the following configuration to enable hashing with only outer IP for GUE packets:

```
Router# configure
Router# hw-module profile gue underlay-hash enable
Router# commit
```

Running Configuration

```
RP/0/RP0/CPU0:ios(config)#show running-config
hw-module profile gue underlay-hash enable
end
```

Verification

Following is the show command output before enabling hashing with only outer IP for GUE packets.

```
RP/0/RP0/CPU0:ios#show dpa objects sys location 0/RP0/CPU0 | include gue
uint32_t gue_ipv4_port => 0
uint32_t gue_ipv6_port => 0
uint32_t gue_mpls_port => 0
ofa_bool_t gue_underlay_hash => FALSE
```

Following is the show command output after enabling hashing with only outer IP for GUE packets.

```
RP/0/RP0/CPU0:ios#show dpa objects sys location 0/RP0/CPU0 | include gue
uint32_t gue_ipv4_port => 0
uint32_t gue_ipv6_port => 0
uint32_t gue_mpls_port => 0
ofa_bool_t gue_underlay_hash => TRUE
```

Flexible Assignment of UDP Port Numbers for Decapsulation

Table 45: Feature History Table

Feature Name	Release Information	Feature Description
Flexible Assignment of UDP Port Numbers for Decapsulation	Release 7.3.3	<p>This feature gives you the flexibility to assign UDP port numbers from 1000 through 6400, through which IPv4, IPv6, and MPLS packets can be decapsulated. Such flexibility allows you to segregate the ingress traffic based on a QoS policy.</p> <p>In earlier releases, you could assign only default ports for decapsulation.</p> <p>The following command is introduced for this feature:</p> <pre>hw-module profile gue udp-dest-port ipv4 <port number> ipv6 <port number> mpls <port number></pre>

This feature provides decapsulation support for GUE packets. In GUE, the payload is encapsulated in an IP packet—IPv4 or IPv6 carrier. The UDP header is added to provide extra hashing parameters and optional payload demultiplexing. At the decapsulation node, the carrier IP and UDP headers are removed, and the packet is forwarded based on the inner payload. Prior to Release 7.3.3, packets were decapsulated using UDP port numbers 6080, 6615, and 6635 for IPv4, IPv6, and MPLS payloads respectively. Starting from Release 7.3.3, you can assign UDP port numbers from 1000 through 64000 to decapsulate IPv4, IPv6, and MPLS packets. Define different port numbers for IPv4, IPv6, and MPLS.

Guidelines for Setting up Decapsulation Using Flexible Port Numbers

Apply these guidelines while assigning flexible port numbers for decapsulation:

Packet	IPv4	IPv6	MPLS
UDP Outer Header	Configure IPv4 port on the hardware module.	Configure IPv6 port on the hardware module.	Configure MPLS port on the hardware module.
Encapsulation Outer Header	Configure an IPv4 encapsulation outer header that matches with the class map source.		
Inner Payload	Note that packets are forwarded based on the inner IPv4 payload.	Note that packets are forwarded based on the inner IPv6 payload.	Note that packets are forwarded based on the inner MPLS payload.

**Note**

- During the decapsulation of the IPv4, IPv6, and MPLS packets, the following headers are removed:
 - The UDP outer header
 - The IPv4 encapsulation outer header
- Select different values for each of these protocols. Valid port numbers are from 1000 through 64000.

Restrictions

The following restrictions are applicable while configuring unique GUE destination port numbers to decapsulate IPv4, IPv6, and MPLS packets using UDP:

- While configuring the tunnel, select one of the following:
 - Match only 16 unique source IP addresses as shown in the example:


```
Router(config-cmap)#match source-address ipv4 210.100.20.0 255.255.255.255
```
 - Match a combination of 64 unique source and destination IP addresses as shown in the example:


```
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
```
- The Classless Inter-Domain Routing (CIDR) value in the source IP address subnet mask must be only /32.
- The destination address subnet mask supports all CIDR values. However, the destination address along with the subnet mask must be unique for all the three UDP payload types—IPv4, IPv6, and MPLS. The configuration fails when the destination IP address and the subnet mask are the same for all three payloads as seen in this example:

```
Router(config)#class-map type traffic match-all SRTE-GUE-DECAP-IPv4
Router(config-cmap)#match destination-address ipv4 10.216.101.0 255.255.255.255
..
Router(config)#class-map type traffic match-all SRTE-GUE-DECAP-IPv6
Router(config-cmap)#match destination-address ipv4 10.216.101.0 255.255.255.255
..
Router(config)#class-map type traffic match-all SRTE-GUE-DECAP-MPLS
Router(config-cmap)#match destination-address ipv4 10.216.101.0 255.255.255.255
..
```

Configuring Port Numbers for Decapsulation

By configuring different port numbers on the destination router, you can match and direct traffic to different paths. For example, traffic for a specific video service can be decapsulated and sent through different ports. The steps that are involved in configuring port numbers for decapsulation are:

1. Configure the UDP destination ports for decapsulation of the required payloads.
2. Configure the traffic class to match the ports.
3. Define a policy map, and associate the traffic class with the traffic policy.

4. Apply the policy for each VRF.



Note For the hardware module flexible port configuration to take effect you must reload the line card.

Configuration Example

```
Hw-module configuration:
=====
Router# configure
Router# hw-module profile gue udp-dest-port ipv4 1001 ipv6 1002 mpls 1003

Class-map configuration:
=====
Router# configure
Router(config)# class-map type traffic match-all udp-v4
Router(config-cmap)# match protocol udp
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match destination-port 1001
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all udp-v6
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match destination-port 1002
Router(config-cmap)# end-class-map
Router(config)# commit

Router(config)# class-map type traffic match-all udp-mpls1
Router(config-cmap)# match protocol udp
Router(config-cmap)# match destination-address ipv4 220.100.20.0 255.255.255.255
Router(config-cmap)# match source-address ipv4 210.100.20.0 255.255.255.255
Router(config-cmap)# match destination-port 1003
Router(config-cmap)# end-class-map
Router(config)# commit

Ingress Policy-map configuration:
=====
Router(config)# policy-map type pbr magic-decap
Router(config-pmap)# class type traffic udp-v4
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-v6
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic udp-mpls1
Router(config-pmap-c)# decapsulate gue variant 1
Router(config-pmap-c)# exit

Router(config-pmap)# class type traffic class-default
Router(config-pmap-c)# exit
Router(config-pmap)# end-policy-map
Router(config)# commit
Router(config)# exit
```

```

Applying policy per VRF:
=====
Router# configure
Router(config)# vrf-policy
Router(config-vrf-policy)# vrf default address-family ipv4 policy type pbr input magic-decap
Router(config-vrf-policy)# commit

```

Running Configuration

```

!! File saved at 16:01:32 UTC Mon Feb 07 2022 by cisco
!! IOS XR Configuration 7.3.3.10I
!! Last configuration change at Mon Feb 7 15:35:11 2022 by cisco
!
logging console disable
username cisco
  group root-lr
  group cisco-support
  secret 10
$6$gHKmELYZAo71BE1.$3KYogrVodJxTRPZgYPGXUXkO4PqQMr2E6oYvJO4ngBmuaGsF2nAB/m1NP5I13zh9HTzBI/k4r8PwWSbsARsmp.
!
vrf vrf-gre
  address-family ipv4 unicast
  !
  address-family ipv6 unicast
  !
!
line console
  exec-timeout 0 0
  absolute-timeout 0
  session-timeout 0
!
line default
  exec-timeout 0 0
  absolute-timeout 0
  session-timeout 0
!
!arp vrf default 29.0.1.2 0000.1122.2929 ARPA
call-home
  service active
  contact smart-licensing
  profile CiscoTAC-1
    active
  destination transport-method http
!
!
ipv6 access-list abf6-gre
  1 permit ipv6 any any nexthop1 ipv6 201:0:1::2
!
ipv4 access-list abf-gre
  1 permit ipv4 any any nexthop1 ipv4 201.0.1.2
!
class-map type traffic match-all udp-v4
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp
  match destination-port 1001
  end-class-map
!
class-map type traffic match-all udp-v6
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp

```

```

    match destination-port 1002
  end-class-map
!
class-map type traffic match-all udp-mpls1
  match destination-address ipv4 220.100.20.0 255.255.255.255
  match source-address ipv4 210.100.20.0 255.255.255.255
  match protocol udp
  match destination-port 1003
end-class-map
!
policy-map type pbr pbr-gre
  class type traffic class-default
    redirect ipv4 nexthop 202.0.1.2
  !
end-policy-map
!
policy-map type pbr magic-decap
  class type traffic udp-v4
    decapsulate gue variant 1
  !
  class type traffic udp-v6
    decapsulate gue variant 1
  !
  class type traffic udp-mpls1
    decapsulate gue variant 1
  !
  class type traffic class-default
  !
end-policy-map
!
interface Bundle-Ether25
  ipv4 address 25.0.1.1 255.255.255.0
  ipv6 address 25:0:1::1/64
  ipv6 enable
  shutdown
!
interface Bundle-Ether28
  ipv4 address 28.0.1.1 255.255.255.0
!
interface Loopback0
  ipv4 address 10.10.10.1 255.255.255.255
!
<output truncated>
interface MgmtEth0/RP0/CPU0/0
  ipv4 address dhcp
!
interface MgmtEth0/RP1/CPU0/0
  ipv4 address dhcp
!
interface BVI23
  ipv4 address 23.0.1.1 255.255.255.0
  ipv6 address 23:0:1::1/64
  ipv6 enable
  shutdown
!
interface BVI29
  ipv4 address 29.0.1.1 255.255.255.0
  ipv6 enable
  shutdown
!
interface HundredGigE0/0/0/0
  shutdown
!

```



```
<output truncated>
l2transport
!
!
interface HundredGigE0/0/0/24
 service-policy type pbr input pbr-gre
 ipv4 address 24.0.1.1 255.255.255.0
 ipv6 address 24:0:1::1/64
 ipv6 enable
!
interface HundredGigE0/0/0/24.24
 ipv4 address 24.0.24.1 255.255.255.0
 ipv6 enable
 encapsulation dot1q 24
!
interface HundredGigE0/0/0/25
 bundle id 25 mode on
!
interface HundredGigE0/0/0/26
 ipv4 address 26.0.1.1 255.255.255.0
 ipv6 address 26:0:1::1/64
 ipv6 enable
!
interface HundredGigE0/0/0/27
 ipv4 address 27.0.1.1 255.255.255.0
 ipv6 enable
!
interface HundredGigE0/0/0/27.27
 ipv4 address 27.0.27.1 255.255.255.0
 ipv6 address 27:0:27::1/64
 ipv6 enable
 shutdown
 encapsulation dot1q 27
!
interface HundredGigE0/0/0/28
 bundle id 28 mode active
!
interface HundredGigE0/0/0/29
 ipv4 address 29.0.1.1 255.255.255.0
 ipv6 enable
!
<output truncated>
interface HundredGigE0/1/0/24
 ipv4 address 124.0.1.1 255.255.255.0
 ipv6 address 124:0:1::1/64
 ipv6 enable
!
<output truncated>
!
interface HundredGigE0/1/0/30
 bundle id 28 mode active
!
interface HundredGigE0/1/0/31
 ipv4 address 31.0.1.1 255.255.255.0
 ipv6 address 31:0:1::1/64
 shutdown
!
<output truncated>
!
route-policy pass
 pass
end-policy
!
router static
```

```

address-family ipv4 unicast
 201.0.1.0/24 tunnel-ip1
 201.0.1.0/24 tunnel-ip2
 201.0.1.0/24 tunnel-ip3
 201.0.1.0/24 tunnel-ip4
!
address-family ipv6 unicast
 201:0:1::/64 tunnel-ip1
 201:0:1::/64 tunnel-ip2
 201:0:1::/64 tunnel-ip3
 201:0:1::/64 tunnel-ip4
!
!
router ospf 10
  router-id 1.1.1.1
  area 0
    ! interface Bundle-Ether28
    interface Loopback0
    !
    interface HundredGigE0/0/0/26
    !
    !
    !
    ! interface HundredGigE0/0/0/27
    ! interface HundredGigE0/0/0/27.27
router bgp 200
  bgp router-id 1.1.1.1
  address-family ipv4 unicast
    maximum-paths ibgp 64
  !
  ! redistribute connected
  ! neighbor 26.0.1.2
  ! remote-as 200
  ! address-family ipv4 unicast
  ! multipath
  ! route-policy pass in
  ! route-policy pass out
  ! next-hop-self
  neighbor 27.0.1.2
  remote-as 200
  address-family ipv4 unicast
    multipath
    route-policy pass in
    route-policy pass out
    next-hop-self
  !
  !
  neighbor 28.0.1.2
  remote-as 200
  address-family ipv4 unicast
    multipath
    route-policy pass in
    route-policy pass out
    next-hop-self
  !
  !
  neighbor 29.0.1.2
  remote-as 200
  address-family ipv4 unicast
    multipath
    route-policy pass in
    route-policy pass out
    next-hop-self
  !

```

```
!
!
vrf-policy
 vrf default address-family ipv4 policy type pbr input magic-decap
!
l2vpn
 bridge group bg
  bridge-domain bd
  ! interface HundredGigE0/0/0/29
  ! static-mac-address 0000.1122.2929
  ! routed interface BVI29
 bridge group bgl
  bridge-domain bdl
  interface HundredGigE0/0/0/23
  static-mac-address 0000.1122.2323
  !
  routed interface BVI23
  !
!
!
!
mpls static
 interface HundredGigE0/0/0/24
 lsp gre
  in-label 35001 allocate per-prefix 202.0.1.2/32
  forward
  path 1 nexthop tunnel-ip1 out-label 35002
  path 2 nexthop tunnel-ip2 out-label 35002
  !
!
!
ssh server vrf default
hw-module profile gue udp-dest-port ipv4 1001 ipv6 1002 mpls 1003
end
```

Verification

Run the **show ofa objects sys location 0/0/CPU0 | inc gue** command in the XR Config mode to verify that the unique GUE port numbers have been configured to decapsulate IPv4, IPv6, and MPLS payloads.

```
Router#show ofa objects sys location 0/0/CPU0 | inc gue
uint32_t gue_ipv4_port => 1001
uint32_t gue_ipv6_port => 1002
uint32_t gue_mpls_port => 1003
```




CHAPTER 15

Configuring 400G Digital Coherent Optics

Table 46: Feature History Table

Feature Name	Release Information	Description
Added Support for DP04QSDD-ER1 and DP01QSDD-ZF1 Optical Modules	Release 24.3.1	<p>This release introduces support for the following Optical Modules on the Cisco 8711-32FH-M router:</p> <ul style="list-style-type: none"> • DP04QSDD-ER1 - Cisco 400G Quad Small Form-Factor Pluggable Double Density (QSFP-DD) • DP01QSDD-ZF1 - Cisco 100G Quad Small Form-Factor Pluggable Double Density (QSFP-DD)
Extended Support for DP04QSDD-HE0 Optical Module	Release 24.1.1	<p>This release introduces support for the Cisco 400G QSFP-DD High-Power (Bright) Optical Module DP04QSDD-HE0, Ethernet Variant on the following routers and line cards -</p> <p>Routers:</p> <ul style="list-style-type: none"> • Cisco 8201 • Cisco 8202 <p>Line cards:</p> <ul style="list-style-type: none"> • 8800-LC-36FH • 88-LC0-36FH-M
Extended Support for DP04QSDD-HE0 Optical Module	Release 7.10.1	<p>This release introduces support for the Cisco 400G QSFP-DD High-Power (Bright) Optical Module DP04QSDD-HE0, Ethernet Variant on the Cisco 8608 router.</p>

Feature Name	Release Information	Description
oFEC Traffic Configuration for QDD-400G-ZRP-S	Release 7.9.1	<p>New Modulation and DAC Rate traffic configurations are supported on QDD-400G-ZRP-S optical module:</p> <ul style="list-style-type: none"> • 400G-TXP-1x1-16 QAM • 4x100G-MXP-1x1-16 QAM • 3x100G-MXP-1x1-8 QAM • 2x100G-MXP-1x1-QPSK • 2x100G-MXP-1x1.25-16 QAM <p>This increases the interoperability of the QDD-400G-ZRP-S optical module across network components supporting these formats.</p>
Support for DP04QSDD-HE0 Optical Module	Release 7.9.1	<p>The Cisco 400G QSFP-DD High-Power (Bright) Optical Module is an enhanced version of the currently available QSFP-DD ZR+ Optical Module. It leverages the same operational modes but provides a major enhancement by increasing the Tx Optical Power up to +1dBm.</p> <p>From this release, the DP04QSDD-HE0 optical module is supported on the Cisco 8201-32FH and Cisco 8201-24H8FH routers.</p>
Support for QDD-400G-ZRP-S Optical Module	Release 7.9.1	<p>This release introduces support for the Cisco 400G QSFP-DD-ZRP-S Ethernet Variant on the Cisco 88-LC0-34H14FH line card.</p>

Cisco offers a range of the new 400G Digital Coherent QSFP-DD optical modules. The optical modules that are available are:

- QDD-400G-ZR-S
- QDD-400G-ZRP-S
- DP04QSDD-HE0
- DP04QSDD-ER1
- DP01QSDD-ZF1

This chapter describes various optical modules and their supported configurations. The following fixed-port routers, line cards, from the indicated Cisco IOS XR software releases, support these optical modules.

Table 47: Fixed-Port Routers and Line Cards that Support various Optical Modules from Indicated Cisco IOS XR Software Releases

Fixed-Port Routers	Optics PID	Minimum IOS XR Software Release
Cisco 8201	QDD-400G-ZR-S	Release 7.3.15
	QDD-400G-ZRP-S	
	DP04QSDD-HE0	Release 24.1.1
Cisco 8202	QDD-400G-ZR-S	Release 7.3.15
	QDD-400G-ZRP-S	
	DP04QSDD-HE0	Release 24.1.1
Cisco 8711-32FH-M	DP04QSDD-ER1	Release 24.3.1
	DP01QSDD-ZF1	
Cisco 8101-32FH	QDD-400G-ZR-S QDD-400G-ZRP-S	Release 7.3.2
Cisco 8201-32FH	DP04QSDD-HE0	Release 7.9.1
Cisco 8201-24H8FH	DP04QSDD-HE0	Release 7.9.1
Cisco 8608	DP04QSDD-HE0	Release 7.10.1
Line Cards	Optics PID	Minimum IOS XR Software Release
8800-LC-36FH	QDD-400G-ZR-S	Release 7.3.15
	QDD-400G-ZRP-S	
	DP04QSDD-HE0	Release 24.1.1
88-LC0-36FH-M	QDD-400G-ZR-S	Release 7.3.15
	QDD-400G-ZRP-S	
	DP04QSDD-HE0	Release 24.1.1
88-LC0-36FH	QDD-400G-ZR-S	Release 7.3.2
	QDD-400G-ZRP-S	
88-LC0-34H14FH	QDD-400G-ZRP-S	Release 7.9.1



Note QDD-400G-ZR-S and QDD-400G-ZRP-S are not supported on 8102-64H fixed-port routers.

QDD-400G-ZRP-S and DP04QSDD-HE0 are not supported on odd-numbered ports of the following routers and line cards:

- Cisco 8201
- Cisco 8202
- 8800-LC-36FH
- 88-LC0-36FH-M

The QDD-400G-ZR-S and QDD-400G-ZRP-S optical modules enable wavelength-division multiplexing (WDM) functionality in the router. These optical modules are DWDM C-band (196.1 THz to 191.3 THz) tunable optical modules. They can be used in both transponder and muxponder modes.

Cisco IOS XR software creates optics and coherent DSP controllers to configure and monitor the performance of the QDD-400G-ZR-S and QDD-400G-ZRP-S optical modules. Optics controllers are used to configure and monitor optical parameters, such as frequency, chromatic dispersion, transmitted output power, modulation, and so on. Coherent DSP controllers are used to monitor network performance parameters like pre- and post-forward error correction (FEC) bit-error rate (pre-FEC BER, post-FEC BER), error corrected bits (EC-BITS), and so on. Forward error correction (FEC) is configured using optical controllers and monitored using coherent DSP controllers.

The QDD-400G-ZR-S and QDD-400G-ZRP-S optical modules support traffic configuration and firmware download. The Cisco IOS XR software collects performance monitoring data and alarms using versatile DOM (VDM).

Due to more power consumption by the QDD-400G-ZR-S and QDD-400G-ZRP-S optical modules, the Cisco IOS XR software operates the fans at a higher speed to cool these optical modules.

The QDD-400G-ZR-S and QDD-400G-ZRP-S optical module configuration is divided into the following categories:

- Traffic configuration – Comprises configuring DAC rate, muxponder mode, modulation, and FEC parameters. Applicable for optics controllers:
 - [Configuring DAC Rate, on page 284](#)
 - [Configuring Muxponder Mode, on page 280](#)
 - [Configuring Modulation, on page 282](#)
 - [Configuring FEC, on page 286](#)
- Optical configuration – Comprises configuring frequency, chromatic dispersion, and optical transmit power. Applicable for optics controllers:
 - [Configuring Frequency, on page 274](#)
 - [Configuring Chromatic Dispersion, on page 276](#)
 - [Configuring Optical Transmit Power, on page 278](#)
- Performance monitoring (PM) – Enables or disables performance monitoring in optical modules. You can also configure PM parameters that comprise signal power, chromatic dispersion, optical signal-to-noise ratio (OSNR), and differential group delay (DGD). Applicable for optics controllers and coherent DSP controllers:
 - [Configuring Performance Monitoring, on page 289](#)

- [Configuring PM Parameters, on page 291](#)
- Loopback configuration – Configures loopback. Applicable for coherent DSP controller:
 - [Configuring Loopback, on page 287](#)
- Alarms threshold configuration – Configures thresholds for monitoring alarms that include optical signal-to-noise ratio (OSNR), differential group delay (DGD), chromatic dispersion (cd high and low), and so on. Applicable for optics controllers:
 - [Configuring Alarms Threshold, on page 294](#)
 - [Configuring FEC Alarm Threshold, on page 297](#)

The following table contains the possible traffic configuration values for the QDD-400G-ZR-S and QDD-400G-ZRP-S optical modules, in the transponder and muxponder mode:

Table 48: 400G Digital Coherent QSFP-DD Traffic Configuration Values

	QDD-400G-ZR-S	QDD-400G-ZRP-S	DP04QSDD-HE0	DP04QSDD-ER1	DP01QSDD-ZF1
Client Speed	1x400G, 4x100G	1x400G, 4x100G, 3x100G, 2x100G, 1x100G Note Release 7.3.15 supports only 1x400 and 4x100 client speed.	1x400G, 4x100G, 3x100G, 2x100G, 1x100G	1x400G, 2x200G, 4x100G	1x100G
Trunk Speed	400G	400G , 300G, 200G, 1x100 Note Release 7.3.15 supports only 400G trunk speed.	400G, 300G, 200G, 100G	400G	100G
Frequency	C-Band, 196.1 To 191.3 THz	C-Band, 196.1 To 191.3 THz	C-Band, 196.1 To 191.3 THz	193.7THz	193.7THz
FEC	cFEC	oFEC, cFEC	oFEC	cFEC, oFEC	oFEC

	QDD-400G-ZR-S	QDD-400G-ZRP-S	DP04QSDD-HE0	DP04QSDD-ER1	DP01QSDD-ZF1
Modulation	16QAM	16QAM, 8QAM, QPSK Release 7.3.15 supports only 16QAM.	16QAM, 8QAM, QPSK	16QAM	QPSK
DAC-Rate	1x1	1x1.25 (oFEC), 1x1 (cFEC)	1x1.25, 1x1	1x1	1x1
Chromatic Dispersion (CD)	-2400 to +2400	Release 7.3.15: -80000 to +80000 Release 7.3.2: -160000 to +160000	-160000 to +160000	-2400 to +2400	-2400 to +2400
Transmitted (Tx) Power	Each optical module has its own transmitting (TX) power range. You can change the transmitting (TX) power value based on the module capability.	Each optical module has its own transmitting (TX) power optimal values. You can change the transmitting (TX) power value based on the module capability.	Each optical module has its own transmitting (TX) power optimal values. You can change the transmitting (TX) power value based on the module capability.	Fixed at maximum output around -9dBm.	Fixed at maximum output around -6dBm.

QDD-400G-ZR-S Transponder and Muxponder Configuration Values

The following table contains the possible Transponder and Muxponder configuration values for the QDD-400G-ZR-S optical module:

Table 49: QDD-400G-ZR-S Transponder and Muxponder Configuration Values

TXP/MXP	Client	Trunk	Modulation	FEC	DAC Rate
400G-TXP	1 client, 400G speed	1 trunk, 400G	16 QAM	cFEC	1x1
4x100G-MXP	4 clients, 100G speed	1 trunk, 400G	16 QAM	cFEC	1x1

QDD-400G-ZRP-S Transponder and Muxponder Configuration Values

The following table contains the possible Transponder and Muxponder configuration values for the QDD-400G-ZRP-S optical module:

Table 50: QDD-400G-ZRP-S Transponder and Muxponder Configuration Values

TXP/MXP	Client	Trunk	Modulation	FEC	DAC Rate	OpenZR+ Support
400G-TXP	1 Client, 400G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1.25	
400G-TXP	1 Client, 400G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1	
400G-TXP	1 Client, 400G speed	1 trunk, 400G speed	16 QAM	cFEC	1x1	
4x100G-MXP	4 clients, 100G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1.25	
4x100G-MXP	4 Client, 100G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1	
4x100G-MXP	4 clients, 100G speed	1 trunk, 400G speed	16 QAM	cFEC	1x1	
3x100G-MXP	3 clients, 100G speed	1 trunk, 400G speed	8 QAM	oFEC	1x1.25	
3x100G-MXP	3 Client, 100G speed	1 trunk, 400G speed	8 QAM	oFEC	1x1	
2x100G-MXP	2 clients, 100G speed	1 trunk, 200G speed	QPSK	oFEC	1x1.50	
2x100G-MXP	2 Client, 100G speed	1 trunk, 400G speed	QPSK	oFEC	1x1	
2x100G-MXP	2 Client, 100G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1.25	
1x100G-MXP	1 client, 100G speed	1 trunk, 100G speed	QPSK	oFEC	1x1.50	

The high optical performance DP04QSDD-HE0 QSFP-DD pluggable coherent optical module is developed for easy deployment in Reconfigurable Optical Add/Drop Multiplexer (ROADM) line systems.

DP04QSDD-HE0 Transponder and Muxponder Configuration Values

The following table contains the possible Transponder and Muxponder configuration values for the DP04QSDD-HE0 optical module:

Table 51: DP04QSDD-HE0 Transponder and Muxponder Configuration Values

TXP/MXP	Client	Trunk	Modulation	FEC	DAC Rate
400G-TXP	1 Client, 400G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1.25

TXP/MXP	Client	Trunk	Modulation	FEC	DAC Rate
100G-TXP	1 Client, 100G speed	1 trunk, 400G speed	QPSK	oFEC	1x1.50
4x100G- MXP	4 clients, 100G speed	1 trunk, 400G speed	16 QAM	oFEC	1x1.25
3x100G-MXP	3 clients, 100G speed	1 trunk, 400G speed	8 QAM	oFEC	1x1.25
2x100-MXP	2 Client, 100G speed	2 Client, 100G speed	QPSK	oFEC	1x1.50

- [Configuring Frequency, on page 274](#)
- [Configuring Chromatic Dispersion, on page 276](#)
- [Configuring Optical Transmit Power, on page 278](#)
- [Configuring Muxponder Mode, on page 280](#)
- [Configuring Modulation, on page 282](#)
- [Configuring DAC Rate, on page 284](#)
- [Configuring FEC, on page 286](#)
- [Configuring Loopback, on page 287](#)
- [Disable Auto-Squelching, on page 288](#)
- [Configuring Performance Monitoring, on page 289](#)
- [Configuring PM Parameters, on page 291](#)
- [Configuring Alarms Threshold, on page 294](#)
- [Configuring FEC Alarm Threshold, on page 297](#)
- [Media Link-down PreFEC Degrade Enablement, on page 302](#)
- [Alarms Troubleshooting, on page 305](#)

Configuring Frequency

You can configure frequency on optics controllers. You can select any C band frequency between the range 196.1 to 191.3 THz, in both ITU and NON-ITU channels.



Note The 100MHz-grid keyword accepts only frequency values as user input. The 50GHz-grid keyword accepts frequency, ITU-channel, or wavelength values as user input. The Cisco IOS XR software then calculates the frequency for a given wavelength or ITU-channel.

Frequency Configuration Example

The following example shows how to configure frequency on the optics controller:

```
Router#config
Router(config)#controller optics 0/2/0/16
Router(config-Optics)#dwdm-carrier 100MHz-grid frequency 1921500
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration:

```
Router#show run controller optics 0/2/0/16
Fri May 28 01:42:32.488 UTC
controller Optics0/2/0/16
  dwdm-carrier 100MHz-grid frequency 1921500
  cd-low-threshold -5000
  cd-high-threshold -5000
!
```

Verification

This example shows how to verify the frequency configuration:

```
Router#show controller optics 0/2/0/16
Fri May 28 01:47:23.953 UTC
Controller State: Up
Transport Admin State: In Service
Laser State: Off
LED State: Off
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSFPDD 400G ZRP
  DWDM carrier Info: C BAND, MSA ITU Channel=80, Frequency=192.15THz,
  Wavelength=1560.200nm
  Alarm Status:
  -----
  Detected Alarms: None
  LOS/LOL/Fault Status:
  Alarm Statistics:
  -----
  HIGH-RX-PWR = 0          LOW-RX-PWR = 0
  HIGH-TX-PWR = 0          LOW-TX-PWR = 0
  HIGH-LBC = 0            HIGH-DGD = 0
  OOR-CD = 0              OSNR = 0
  WVl-OOL = 0             MEA = 0
  IMPROPER-REM = 0
  TX-POWER-PROV-MISMATCH = 0
  Laser Bias Current = 0.0 mA
  Actual TX Power = -40.00 dBm
  RX Power = -40.00 dBm
  RX Signal Power = -40.00 dBm
  Frequency Offset = 0 MHz
  Laser Temperature = 0.00 Celsius
  Laser Age = 0 %
  DAC Rate = 1x1.25
  Performance Monitoring: Enable
  THRESHOLD VALUES
  -----
  Parameter                High Alarm  Low Alarm  High Warning  Low Warning
  -----
  Rx Power Threshold(dBm)   13.0       -24.0      10.0          -22.0
  Tx Power Threshold(dBm)   0.0        -16.0      -2.0          -14.0
  LBC Threshold(mA)         0.00       0.00      0.00          0.00
  Temp. Threshold(celsius)  80.00     -5.00     75.00         0.00
  Voltage Threshold(volt)   3.46       3.13      3.43          3.16
  LBC High Threshold = 98 %
  Configured Tx Power = -10.00 dBm
  Configured CD High Threshold = -5000 ps/nm
  Configured CD lower Threshold = -5000 ps/nm
  Configured OSNR lower Threshold = 9.00 dB
  Configured DGD Higher Threshold = 80.00 ps
  Baud Rate = 60.1385459900 GBd
```

```

Modulation Type: 16QAM
Chromatic Dispersion 0 ps/nm
Configured CD-MIN -26000 ps/nm CD-MAX 26000 ps/nm
Second Order Polarization Mode Dispersion = 0.00 ps^2
Optical Signal to Noise Ratio = 0.00 dB
Polarization Dependent Loss = 0.00 dB
Polarization Change Rate = 0.00 rad/s
Differential Group Delay = 0.00 ps
Temperature = 21.00 Celsius
Voltage = 3.42 V
Transceiver Vendor Details
  Form Factor           : QSFP-DD
  Optics type          : QSFPDD 400G ZRP
  Name                 : CISCO-ACACIA
  OUI Number           : 7c.b2.5c
  Part Number          : DP04QSDD-E30-19E
  Rev Number           : 10
  Serial Number        : ACA244900GN
  PID                  : QDD-400G-ZRP-S
  VID                  : ES03
  Firmware Version     : 161.06
  Date Code (yy/mm/dd) : 20/12/08
!

```

Configuring Chromatic Dispersion

You can configure chromatic dispersion on optics controllers. When you configure the maximum and minimum values for chromatic dispersion for any data rate, ensure that the minimum difference between the configured values is equal to or greater than 1000 ps/nm.

The following table lists the default CD search range:

Table 52: Default CD Search Range

Muxponder Rate	FEC Value	Default CD Search Range (Min-Max)
400	OFEC	-26000 to +26000
400	CFEC	-2400 to +2400
300	OFEC	-50000 to +50000
200	OFEC	-50000 to +50000
100	OFEC	-80000 to +80000

Chromatic Dispersion Configuration Example

This example shows how to configure chromatic dispersion on the optics controller:

```

Router#configure
Router(config)#controller optics 0/0/0/13
Router(config-Optics)#cd-max 4000
Router(config-Optics)#cd-min -4000
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit

```

Running Configuration

This example shows the running configuration for the optics controller:

```
Router#show run controller optics 0/0/0/13
Thu May 13 12:24:42.353 UTC
controller Optics0/0/0/13
  cd-min -4000
  cd-max 4000
!
```

Verification

This example shows how to verify the configured chromatic dispersion values for the optics controller:

```
Router#show controller optics 0/0/0/13
Controller State: Up
Transport Admin State: In Service
Laser State: On
LED State: Green
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSPDD 400G ZR
  DWDM carrier Info: C BAND, MSA ITU Channel=61, Frequency=193.10THz,
  Wavelength=1552.524nm
  Alarm Status:
  -----
  Detected Alarms: None
  LOS/LOL/Fault Status:
  Alarm Statistics:
  -----
  HIGH-RX-PWR = 0          LOW-RX-PWR = 0
  HIGH-TX-PWR = 0          LOW-TX-PWR = 0
  HIGH-LBC = 0            HIGH-DGD = 0
  OOR-CD = 0              OSNR = 35
  WVL-OOL = 0             MEA = 0
  IMPROPER-REM = 0
  TX-POWER-PROV-MISMATCH = 0
  Laser Bias Current = 0.0 %
  Actual TX Power = -7.87 dBm
  RX Power = -8.27 dBm
  RX Signal Power = -8.43 dBm
  Frequency Offset = 130 MHz
  Performance Monitoring: Enable
  THRESHOLD VALUES
  -----
  Parameter                High Alarm  Low Alarm  High Warning  Low Warning
  -----
  Rx Power Threshold(dBm)   1.9        -28.2      0.0           -25.0
  Tx Power Threshold(dBm)   0.0        -15.0      -2.0          -16.0
  LBC Threshold(mA)         0.00       0.00      0.00          0.00
  Temp. Threshold(celsius)  80.00     -5.00     75.00         15.00
  Voltage Threshold(volt)   3.46       3.13      3.43          3.16
  LBC High Threshold = 98 %
  Configured Tx Power = -6.00 dBm
  Configured CD High Threshold = 80000 ps/nm
  Configured CD lower Threshold = -80000 ps/nm
  Configured OSNR lower Threshold = 9.00 dB
  Configured DGD Higher Threshold = 80.00 ps
  Baud Rate = 59.8437500000 GBd
  Modulation Type: 16QAM
  Chromatic Dispersion 0 ps/nm
  Configured CD-MIN -4000 ps/nm CD-MAX 4000 ps/nm
  Second Order Polarization Mode Dispersion = 5.00 ps^2
  Optical Signal to Noise Ratio = 36.30 dB
```

```

Polarization Dependent Loss = 0.40 dB
Polarization Change Rate = 0.00 rad/s
Differential Group Delay = 4.00 ps
Temperature = 54.00 Celsius
Voltage = 3.37 V
Transceiver Vendor Details
Form Factor           : QSFP-DD
Optics type          : QSFPDD 400G ZR
Name                 : CISCO-ACACIA
OUI Number          : 7c.b2.5c
Part Number         : DP04QSDD-E20-19E
Rev Number          : 10
Serial Number       : ACA2447003L
PID                 : QDD-400G-ZR-S
VID                 : ES03
Firmware Version    : 61.12
Date Code (yy/mm/dd) : 20/12/02

```

Configuring Optical Transmit Power

You can set the transmit power of the optical signal.

Each QDD-400G-ZR-S and QDD-400G-ZRP-S optical module has its own optical transmit (TX) power range. You can change the optical transmit (TX) power value based on the module capability. For "Transmitter specifications", see the [Cisco 400G Digital Coherent Optics QSFP-DD Optical Modules Data Sheet](#).

Table 53: Optical Transmit Power Values

Optical Module	Trunk Speed ^{1,3}	Optical Transmit Power (Tx) Shaping	Interval	Supported Range of Optical Transmit Power (Tx) Values (in units of 0.1dBm) ²		
				Minimum Value	Maximum Value - Typical	Maximum Value - Worst Case
QDD-400G-ZR-S	400G	No	1	-150	-100	-100
QDD-400G-ZRP-S	400G	Yes	1	-150	-110	-130
	300G			-150	-104	-119
	200G			-150	-90	-105
	100G			-150	-59	-75
DP04QSDD-HE0	400G	Yes	1	-100	20	10
	300G					
	200G					
	100G					

¹. Release 7.3.15 supports 4x100G muxponder mode or trunk speed.

². The default optical transmit power (Tx) value is -10 dBm, however with Tx shaping enabled the maximum power in 1x400G, 4x100G, 3x100G, 2x100G, and 1x100G modes may be less than -10 dBm.

³. Release 7.3.2 and future releases support 3x100G, 2x100G, and 1x100G muxponder modes or trunk speed.

Transmitting Power Configuration Example

The following example shows how to configure the optical transmit (TX) power on the optics controller:

```
Router#config
Router(config)#controller optics 0/2/0/16
Router(config-Optics)#transmit-power -125
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration for the optics controller:

```
Router#show run controller optics 0/2/0/16
Thu May 13 12:52:35.020 UTC
controller Optics0/0/0/1
  cd-min -4000
  cd-max 4000
  transmit-power -125
!
```

Verification

This example shows how to verify the configured optical transmit power for the optics controller:

```
Router#show controller optics 0/2/0/16
Fri May 28 02:52:06.182 UTC
Controller State: Up
Transport Admin State: In Service
Laser State: Off
LED State: Off
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSFPDD 400G ZRP
  DWDM carrier Info: C BAND, MSA ITU Channel=80, Frequency=192.15THz,
  Wavelength=1560.200nm
  Alarm Status:
  -----
  Detected Alarms: None
  LOS/LOL/Fault Status:
  Alarm Statistics:
  -----
  HIGH-RX-PWR = 0          LOW-RX-PWR = 0
  HIGH-TX-PWR = 0          LOW-TX-PWR = 0
  HIGH-LBC = 0            HIGH-DGD = 0
  OOR-CD = 0              OSNR = 0
  WVL-OOL = 0             MEA = 0
  IMPROPER-REM = 0
  TX-POWER-PROV-MISMATCH = 0
  Laser Bias Current = 0.0 mA
  Actual TX Power = -40.00 dBm
  RX Power = -40.00 dBm
  RX Signal Power = -40.00 dBm
  Frequency Offset = 0 MHz
  Laser Temperature = 0.00 Celsius
  Laser Age = 0 %
  DAC Rate = 1x1.25
  Performance Monitoring: Enable
  THRESHOLD VALUES
  -----
  Parameter                High Alarm  Low Alarm  High Warning  Low Warning
  -----
```

```

Rx Power Threshold(dBm)      13.0      -24.0      10.0      -22.0
Tx Power Threshold(dBm)      0.0      -16.0      -2.0      -14.0
LBC Threshold(mA)            0.00     0.00     0.00     0.00
Temp. Threshold(celsius)     80.00    -5.00    75.00    0.00
Voltage Threshold(volt)      3.46     3.13     3.43     3.16
LBC High Threshold = 98 %

```

Configured Tx Power = -12.50 dBm

```

Configured CD High Threshold = -5000 ps/nm
Configured CD lower Threshold = -5000 ps/nm
Configured OSNR lower Threshold = 9.00 dB
Configured DGD Higher Threshold = 80.00 ps
Baud Rate = 60.1385459900 GBd
Modulation Type: 16QAM
Chromatic Dispersion 0 ps/nm
Configured CD-MIN -4000 ps/nm CD-MAX 4000 ps/nm
Second Order Polarization Mode Dispersion = 0.00 ps^2
Optical Signal to Noise Ratio = 0.00 dB
Polarization Dependent Loss = 0.00 dB
Polarization Change Rate = 0.00 rad/s
Differential Group Delay = 0.00 ps
Temperature = 20.00 Celsius
Voltage = 3.41 V

```

Transceiver Vendor Details

```

Form Factor      : QSFP-DD
Optics type     : QSFPDD 400G ZRP
Name            : CISCO-ACACIA
OUI Number      : 7c.b2.5c
Part Number     : DP04QSDD-E30-19E
Rev Number      : 10
Serial Number   : ACA244900GN
PID             : QDD-400G-ZRP-S
VID             : ES03
Firmware Version : 161.06
Date Code(yy/mm/dd) : 20/12/08

```

Configuring Muxponder Mode

By default, the Cisco IOS XR software configures the QDD-400G-ZR-S and QDD-400G-ZRP-S optical modules in the 400G transponder mode.

Using the **breakout muxponder mode** command, you can configure muxponder mode on optics controllers. Based on the muxponder mode, you can choose the modulation.

Muxponder mode options available for QDD-400G-ZR-S are:

- 4x100

Muxponder mode options available for QDD-400G-ZRP-S are:

- 4x100
- 3x100
- 2x100



Note Release 7.3.15 supports only 4x100 muxponder mode.

See the following tables for the modulation values, based on the muxponder mode:

- [QDD-400G-ZR-S Transponder and Muxponder Configuration Values, on page 272](#)
- [QDD-400G-ZRP-S Transponder and Muxponder Configuration Values, on page 272](#)

Using the **no breakout muxponder mode** command, you can switch from the muxponder mode to the transponder mode, on optics controllers.

Muxponder Mode Configuration Example

The following example shows how to configure muxponder mode on the optics controller:

```
Router#config
Router(config)#controller optics 0/0/0/13
Router(config-Optics)#breakout 4x100
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit
```



Note In the above example, the Cisco IOS XR software creates four Ethernet clients with 100GE speed, which can be verified using the **show interfaces brief | include R/S/I/P** command.

Running Configuration

This example shows the running configuration for the optics controller:

```
Router#show run controller optics 0/0/0/13
Thu May 13 12:24:42.353 UTC
controller Optics0/0/0/13
  cd-min -4000
  cd-max 4000
  breakout 4x100
  !
```

Verification

This example shows how to verify the muxponder mode configuration:

```
Router#show interfaces brief | include 0/0/0/13
Hu0/0/0/13/0      up      up      ARPA  1514  100000000
Hu0/0/0/13/1      up      up      ARPA  1514  100000000
Hu0/0/0/13/2      up      up      ARPA  1514  100000000
Hu0/0/0/13/3      up      up      ARPA  1514  100000000
```

Transponder Mode Configuration Example

The following example shows how to switch to the transponder mode, on the optics controller:

```
Router#config
Router(config)#controller optics 0/0/0/13
Router(config-Optics)#no breakout 4x100
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit
```



Note The Cisco IOS XR software creates a single 400GE interface, which can be verified using the **show interfaces brief | include R/S/I/P** command.

Running Configuration

This example shows the running configuration for the optics controller. The breakout configuration is absent in the running configuration.

```
Router#show run controller optics 0/0/0/13
Thu May 13 13:51:20.330 UTC
controller Optics0/0/0/13
  cd-min -4000
  cd-max 4000
  transmit-power -100
!
```

Verification

This example shows how to verify the transponder mode configuration:

```
Router#show interfaces brief | include 0/0/0/13
FH0/0/0/13          up          up          ARPA  1514  400000000
```

Configuring Modulation

You can configure modulation on optics controllers. Based on the muxponder mode, you can choose the modulation.



Note The system accepts any modulation value that is entered. However, if the modulation value is outside the supported range, it is not configured on the optical module. Instead, the optical module is auto-configured with a valid modulation value. To view this value, use the **show controller optics R/S/I/P** command.

See the following tables for the supported modulation values:

- [QDD-400G-ZR-S Transponder and Muxponder Configuration Values, on page 272](#)
- [QDD-400G-ZRP-S Transponder and Muxponder Configuration Values, on page 272](#)

Modulation Configuration Example

The following example shows how to configure modulation on the optics controller:

```
Router#config
Router(config)#controller optics 0/0/0/1
Router(config-Optics)#modulation 16Qam
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration:

```
Router#show run controller optics 0/0/0/1
controller Optics0/0/0/1
  cd-min -4000
  cd-max 4000
  transmit-power -100
  modulation 16Qam
!
```



Note Use the `show controller optics R/S/I/P` command to verify the modulation value of the optical module.

Verification

This example shows how to verify the configured modulation value for the optics controller:

```

Router#show controller optics 0/0/0/1
Controller State: Up
Transport Admin State: In Service
Laser State: On
LED State: Green
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSFPDD 400G ZR
  DWDM carrier Info: C BAND, MSA ITU Channel=61, Frequency=193.10THz,
  Wavelength=1552.524nm
  Alarm Status:
  -----
  Detected Alarms: None
  LOS/LOL/Fault Status:
  Alarm Statistics:
  -----
  HIGH-RX-PWR = 0           LOW-RX-PWR = 0
  HIGH-TX-PWR = 0           LOW-TX-PWR = 0
  HIGH-LBC = 0              HIGH-DGD = 0
  OOR-CD = 0                OSNR = 35
  WV-L-OOL = 0              MEA = 0
  IMPROPER-REM = 0
  TX-POWER-PROV-MISMATCH = 0
  Laser Bias Current = 0.0 %
  Actual TX Power = -7.87 dBm
  RX Power = -8.27 dBm
  RX Signal Power = -8.43 dBm
  Frequency Offset = 130 MHz
  Performance Monitoring: Enable
  THRESHOLD VALUES
  -----
  Parameter                High Alarm  Low Alarm  High Warning  Low Warning
  -----
  Rx Power Threshold(dBm)   1.9         -28.2      0.0           -25.0
  Tx Power Threshold(dBm)   0.0         -15.0      -2.0          -16.0
  LBC Threshold(mA)         0.00        0.00       0.00          0.00
  Temp. Threshold(celsius)  80.00       -5.00      75.00         15.00
  Voltage Threshold(volt)   3.46        3.13       3.43          3.16
  LBC High Threshold = 98 %
  Configured Tx Power = -6.00 dBm
  Configured CD High Threshold = 80000 ps/nm
  Configured CD lower Threshold = -80000 ps/nm
  Configured OSNR lower Threshold = 9.00 dB
  Configured DGD Higher Threshold = 80.00 ps
  Baud Rate = 59.8437500000 GBd
Modulation Type: 16QAM
  Chromatic Dispersion 0 ps/nm
  Configured CD-MIN -4000 ps/nm CD-MAX 4000 ps/nm
  Second Order Polarization Mode Dispersion = 5.00 ps^2
  Optical Signal to Noise Ratio = 36.30 dB
  Polarization Dependent Loss = 0.40 dB
  Polarization Change Rate = 0.00 rad/s
  Differential Group Delay = 4.00 ps
  Temperature = 54.00 Celsius
  Voltage = 3.37 V
    
```

```

Transceiver Vendor Details
  Form Factor           : QSFP-DD
  Optics type          : QSFPDD 400G ZR
  Name                 : CISCO-ACACIA
  OUI Number           : 7c.b2.5c
  Part Number          : DP04QSDD-E20-19E
  Rev Number           : 10
  Serial Number        : ACA2447003L
  PID                  : QDD-400G-ZR-S
  VID                  : ES03
  Firmware Version     : 61.12
  Date Code (yy/mm/dd) : 20/12/02

```

Configuring DAC Rate

You can set the DAC (digital to analog conversion) sampling rate on optics controllers. You can modify the DAC sampling rate only on the QDD-400G-ZRP-S and DP04QSDD-HE optical module.



Note QDD-400G-ZR-S supports 1x1 dac-rate in cFEC mode. QDD-400G-ZRP-S and DP04QSDD-HE supports 1x1 dac-rate in cFEC mode and 1x1.25 dac-rate in oFEC mode.

DAC Rate Configuration Example

The following example shows how to set the DAC rate on the optics controller:

```

Router#config
Router(config)#controller optics 0/0/0/1
Router(config-Optics)#dac-rate 1x1

```

Verification

This example shows the running configuration:

```

Router#show run controller optics 0/0/0/1
Thu May 13 12:52:35.020 UTC
controller Optics0/0/0/1
  cd-min -4000
  cd-max 4000
  transmit-power -100
  modulation 16Qam
  DAC-Rate 1x1
!
!

```

Verification

This example shows how to verify the configured DAC rate for the optics controller:

```

Router#show controller optics 0/0/0/1
Controller State: Up
Transport Admin State: In Service
Laser State: On
LED State: Green
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSFPDD 400G ZR
  DWDM carrier Info: C BAND, MSA ITU Channel=61, Frequency=193.10THz,
  Wavelength=1552.524nm
  Alarm Status:

```

```

-----
Detected Alarms: None
LOS/LOL/Fault Status:
Alarm Statistics:
-----
HIGH-RX-PWR = 0          LOW-RX-PWR = 0
HIGH-TX-PWR = 0          LOW-TX-PWR = 0
HIGH-LBC = 0            HIGH-DGD = 0
OOR-CD = 0              OSNR = 35
WVL-OOL = 0            MEA = 0
IMPROPER-REM = 0
TX-POWER-PROV-MISMATCH = 0
Laser Bias Current = 0.0 %
Actual TX Power = -7.87 dBm
RX Power = -8.27 dBm
RX Signal Power = -8.43 dBm
Frequency Offset = 130 MHz
DAC Rate = 1x1
Performance Monitoring: Enable
THRESHOLD VALUES
-----
Parameter                High Alarm  Low Alarm  High Warning  Low Warning
-----
Rx Power Threshold(dBm)   1.9        -28.2     0.0           -25.0
Tx Power Threshold(dBm)   0.0        -15.0     -2.0          -16.0
LBC Threshold(mA)        0.00       0.00     0.00          0.00
Temp. Threshold(celsius) 80.00      -5.00     75.00         15.00
Voltage Threshold(volt)   3.46       3.13     3.43          3.16
LBC High Threshold = 98 %
Configured Tx Power = -6.00 dBm
Configured CD High Threshold = 80000 ps/nm
Configured CD lower Threshold = -80000 ps/nm
Configured OSNR lower Threshold = 9.00 dB
Configured DGD Higher Threshold = 80.00 ps
Baud Rate = 59.8437500000 GBd
Modulation Type: 16QAM
Chromatic Dispersion 0 ps/nm
Configured CD-MIN -4000 ps/nm  CD-MAX 4000 ps/nm
Second Order Polarization Mode Dispersion = 5.00 ps^2
Optical Signal to Noise Ratio = 36.30 dB
Polarization Dependent Loss = 0.40 dB
Polarization Change Rate = 0.00 rad/s
Differential Group Delay = 4.00 ps
Temperature = 54.00 Celsius
Voltage = 3.37 V
Transceiver Vendor Details
Form Factor                : QSFP-DD
Optics type                : QSFPDD 400G ZR
Name                      : CISCO-ACACIA
OUI Number                 : 7c.b2.5c
Part Number                : DP04QSDD-E20-19E
Rev Number                 : 10
Serial Number              : ACA2447003L
PID                       : QDD-400G-ZR-S
VID                       : ES03
Firmware Version           : 61.12
Date Code(yy/mm/dd)       : 20/12/02

```

Configuring FEC

You can configure forward error correction (FEC) only on optics controllers. You can modify FEC only on the QDD-400G-ZRP-S and DP04QSDD-HE optical module. FEC is a feature that is used for controlling errors during data transmission. This feature works by adding data redundancy to the transmitted message using an algorithm. This redundancy allows the receiver to detect and correct a limited number of errors occurring anywhere in the message, instead of having to ask the transmitter to resend the message.



Note QDD-400G-ZR-S supports cFEC (concatenated forward error correction). QDD-400G-ZRP-S and DP04QSDD-HE supports cFEC and oFEC (open forward error correction).

FEC Configuration Example

The following sample shows how to configure FEC on the optics controller:

```
Router#configure
Router(config)#controller optics 0/0/0/13
Router(config-Optics)#fec CFEC
Router(config-Optics)#commit
Router(config-Optics)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration:

```
Router#show controllers optics 0/0/0/13
controller Optics0/0/0/1
  cd-min -4000
  cd-max 4000
  transmit-power -100
  fec CFEC
  modulation 16Qam
  DAC-Rate 1x1.25
!
```

Verification

This example shows how to verify the FEC configuration for the optics controller:

```
Router#show controller coherentdsp 0/0/0/13
Thu May 27 17:28:51.960 UTC
Port                : CoherentDSP 0/0/0/13
Controller State    : Down
Inherited Secondary State : Normal
Configured Secondary State : Maintenance
Derived State       : Maintenance
Loopback mode       : Internal
BER Thresholds      : SF = 1.0E-5  SD = 1.0E-7
Performance Monitoring : Enable
Bandwidth           : 400.0Gb/s

Alarm Information:
LOS = 6 LOF = 0 LOM = 0
OOF = 0 OOM = 0 AIS = 0
IAE = 0 BIAE = 0      SF_BER = 0
SD_BER = 0      BDI = 0 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0      FLEXO_GIDM = 0
```



```

FLEXO-MM = 0      FLEXO-LOM = 0      FLEXO-RDI = 0
FLEXO-LOF = 5
Detected Alarms                                     : LOS
Bit Error Rate Information
PREFEC BER                                          : 5.0E-01
POSTFEC BER                                        : 0.0E+00
Q-Factor                                           : 0.00 dB
Q-Margin                                           : -7.20dB
OTU TTI Received

FEC mode                                           : C_FEC

```

Configuring Loopback

You can configure internal or line loopback on coherent DSP controllers. Loopback can be performed only in the maintenance mode.



Note Line loopback mode is supported only on Cisco 8000 series line cards and fixed-port routers based on Q100 and Q200 silicon.

Loopback Configuration Example

This example shows how to enable internal loopback configuration on coherent DSP controllers:

```

Router#config
Router(config)#controller coherentDSP 0/0/0/4
Router(config-CoDSP)#secondary-admin-state maintenance
Router(config-CoDSP)#loopback internal
Router(config-CoDSP)#commit

```

Running Configuration

This example shows the running configuration on coherent DSP controllers:

```

Router#show run controller coherentdsp 0/0/0/4
Thu May 13 19:51:08.175 UTC
controller CoherentDSP0/0/0/4
  secondary-admin-state maintenance
  loopback internal
!
```

Verification

This example shows how to verify the loopback configuration on coherent DSP controllers:

```

Router#show controller coherentdsp 0/0/0/4
Thu May 27 17:28:51.960 UTC
Port                                     : CoherentDSP 0/0/0/4
Controller State                         : Down
Inherited Secondary State               : Normal
Configured Secondary State           : Maintenance
Derived State                           : Maintenance
Loopback mode                       : Internal
BER Thresholds                          : SF = 1.0E-5   SD = 1.0E-7
Performance Monitoring                  : Enable
Bandwidth                               : 400.0Gb/s
Alarm Information:
LOS = 6  LOF = 0  LOM = 0
OOF = 0  OOM = 0  AIS = 0

```

```

IAE = 0 BIAE = 0          SF_BER = 0
SD_BER = 0          BDI = 0 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0          FLEXP_GIDM = 0
FLEXP-MM = 0          FLEXP-LOM = 0          FLEXP-RDI = 0
FLEXP-LOF = 5
Detected Alarms                               : LOS
Bit Error Rate Information
PREFEC BER                                     : 5.0E-01
POSTFEC BER                                    : 0.0E+00
Q-Factor                                       : 0.00 dB
Q-Margin                                       : -7.20dB
OTU TTI Received
FEC mode                                       : C_FEC

```

Disable Auto-Squelching

Table 54: Feature History Table

Feature Name	Release Information	Description
Disable Auto-Squelching	Release 7.11.1	<p>This release introduces support to disable Auto squelching. This helps to detect weak signals that are hidden within the laser source noise. By disabling Auto squelch, you can reduce the processing overhead in systems that have stable laser sources and minimal noise, helping you optimize the performance of your system. When the Auto squelch function is enabled, the optical module will generate a local fault signal on the host side if it detects a fault on the media side. By default, Auto squelch is enabled.</p> <p>The feature introduces these changes:</p> <p>CLI:The following keyword has been added.</p> <ul style="list-style-type: none"> • host auto-squelch disable <p>YANG DATA models:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-controller-optics-cfg</code> (see Github, YANG Data Models Navigator)

This release introduces the support to disable auto-squelch functionality on the module on the host side. When enabled, the squelch function is activated on the module when no suitable media-side input signal from the

remote end is available to be forwarded to the host-side output (example: Rx LOS is asserted). Auto squelching is commonly used to suppress unwanted noise from laser sources in communication systems. When disabled and no valid signal is detected on the module from the remote end, the module will generate a local fault towards the NPU. However, disabling auto-squelching provides you with expanded signal detection. This enables you to detect extremely weak signals that are embedded within the laser source noise. Also, by eliminating the need to continuously monitor and suppress unwanted noise, system resources can be allocated more efficiently, leading to improved performance.

In this feature, we introduced the **host auto-squelch disable** command to disable the auto-squelch functionality when there is an invalid input signal from the remote end. This feature provides you with the flexibility to customize the system's behavior according to your requirements.

Disabling Laser Squelching Configuration Example

This example shows how to disable laser squelching for a host on controller optics:

```
router#config
router(config)#controller 0/0/0/0
router(config-Optics)#host auto-squelch disable
router(config-Optics)#commit
```

Verification

This example shows how to verify the laser squelching disabled configuration:

```
router#show controllers optics 0/0/0/0
Host Squelch Status: disable
```

Configuring Performance Monitoring

Feature Name	Release Information	Description
New Parameters for Performance Monitoring	Release 24.2.11	<p>In this release, we've added two new parameters for performance monitoring. These parameters allow you to check the quality of electrical signals between an ASIC or NPU and the 400G Digital Coherent QSFP-DD optical modules periodically. This helps to detect errors occurring during data transmission and link initialization.</p> <p>This feature introduces new parameters Host-Intf-{n}-FEC-BER and Host-Intf-{n}-FEC-FERC in the show controllers coherentdsp command.</p>

Performance monitoring (PM) parameters are used by service providers to gather, store, set thresholds for, and report performance data for early detection of problems. The user can retrieve both current and historical PM counters for the various controllers in 30-second, 15-minute, and 24-hour intervals.

Performance monitoring can be configured on optics controllers and coherent DSP controllers.

To stop performance monitoring on optics or coherent DSP controllers, use the **perf-mon disable** keyword.

Release 24.2.11 introduces the following host PM parameters:

- **Host-Intf-{n}-FEC-BER**
- **Host-Intf-{n}-FEC-FERC**

The following table lists the details of the host PM parameters:

Table 55: Performance Monitoring Parameters

PM Parameter	Mode Type	Number of Host Interfaces	Description
Host-Intf-{n}-FEC-BER	Transponder	1	$n=0$ For example, Host-Intf-0-FEC-BER
	Muxponder	4	$n = 0, 1, 2, \text{ and } 3.$ For example: <ul style="list-style-type: none"> • Host-Intf-0-FEC-BER • Host-Intf-1-FEC-BER • Host-Intf-2-FEC-BER • Host-Intf-3-FEC-BER
Host-Intf-{n}-FEC-FERC	Transponder	1	$n=0$ For example, Host-Intf-0-FEC-FERC
	Muxponder	4	$n = 0, 1, 2, \text{ and } 3$ For example: <ul style="list-style-type: none"> • Host-Intf-0-FEC-FERC • Host-Intf-1-FEC-FERC • Host-Intf-2-FEC-FERC • Host-Intf-3-FEC-FERC

These parameters check the quality of electrical signals between an ASIC or NPU and the 400G Digital Coherent QSFP-DD optical modules. The Cisco IOS XR software installed on the router collects Bit Error Rate (BER) and Frame Error Counters (FERC) information from the optical modules and maintains history for 30-sec, 15-minute, and 24-hours intervals.



Note You can check these signals for host interfaces in both transponder and muxponder mode.

Configuring PM Parameters

The performance monitoring (PM) threshold and the threshold crossing alert (TCA) reporting status can be configured for optics controllers and coherent DSP controllers:

Table 56: PM Thresholds and TCA Report Status for Optics Controllers

PM Parameters	Description
CD	Sets the CD (chromatic dispersion) threshold or TCA reporting status.
DGD	Sets the DGD (differential group delay) threshold or TCA reporting status.
LBC	Sets the LBC (laser bias current) threshold or TCA reporting status in mA.
FREQ-OFF	Sets the FREQ-OFF (low signal frequency offset) threshold or TCA reporting status in Mhz.
OPR	Sets the OPR (optical power RX) threshold or TCA reporting status in uW or dbm.
OPT	Sets the OPT (optical power TX) threshold or TCA reporting status in uW or dbm.
OSNR	Sets the OSNR (optical signal-to-noise ratio) threshold or TCA reporting status.
PCR	Sets the PCR (polarization change rate) threshold or TCA reporting status.
PDL	Sets the PDL (polarization dependent loss) threshold or TCA reporting status.
RX-SIG	Sets the RX-SIG (receiving signal power) threshold or TCA reporting status in uW or dbm.
SNR	Sets the SNR (signal-to-noise ratio) threshold or TCA reporting status.
SOPMD	Sets the SOPMD (second order polarization mode dispersion) threshold or TCA reporting status.

Table 57: PM Thresholds TCA Report Status for Coherent DSP Controllers

PM Parameters	Description
Q	Sets the Q threshold or TCA reporting status.
Q-margin	Sets the Q margin threshold or TCA reporting status.
EC-BITS	Sets the EC-BITS (error corrected bits) threshold or TCA reporting status.
PostFEC BER	Sets the post-FEC BER threshold or TCA reporting status.
PreFEC BER	Sets the pre-FEC BER threshold or TCA reporting status.
UC-WORDS	Sets the UC-WORDS (uncorrected words) threshold or TCA reporting status.
Host-Intf-0-FEC-BER	<p>Sets the Host-Intf-0-FEC-BER threshold or TCA reporting status, where:</p> <ul style="list-style-type: none"> • AVG - specifies the number of corrected bits received from the host interface prior to a PM interval. • MIN - specifies the minimum number of corrected bits received from the host interface over a sub-interval and prior to a PM interval. • MAX - specifies the maximum number of corrected bits received from the host interface over a sub-interval and prior to a PM interval.
Host-Intf-0-FEC-FERC	<p>Sets the Host-Intf-0-FEC-FERC threshold or TCA reporting status, where:</p> <ul style="list-style-type: none"> • AVG - specifies the number of frames received from the host interface during a sub-interval. • MIN - specifies the minimum number of frames received from the host interface with uncorrected errors over a sub-interval and prior to a PM interval. • MAX - specifies the maximum number of frames received from the host interface with uncorrected errors over a sub-interval and prior to a PM interval.

Performance Monitoring Configuration Example

This example shows how to enable performance monitoring and set PM thresholds on the optics controller:

```

Router#config
Router(config)#controller optics 0/2/0/16
Router(config-Optics)#perf-mon enable
Router(config-Optics)#pm 30-sec optics threshold cd max 100
Router(config-Optics)#pm 30-sec optics threshold cd min -100
Router(config-Optics)#commit

```

Running Configuration

This example shows the running configuration on optics controllers:

```

Router#show run controller optics 0/2/0/16
Thu May 13 20:18:55.957 UTC
controller Optics0/2/0/16
pm 30-sec optics threshold cd max 100
pm 30-sec optics threshold cd min -100
perf-mon enable
!

```

Verification

This example shows how to verify the PM parameters on optics controllers. Verify the configuration changes in the Configured Threshold fields:

```

Router#show controller optics 0/2/0/16 pm current 30-sec optics 1
Thu May 27 17:58:49.889 UTC
Optics in the current interval [17:58:30 - 17:58:49 Thu May 27 2021]
Optics current bucket type : Valid

```

	MIN	AVG	MAX	Operational	Configured	TCA	Operational
	Configured	TCA		Threshold(min)	Threshold(min)	(min)	Threshold(max)
LBC[mA]	: 0.0	0.0	0.0	0.0	NA	NO	100.0
	NA	NO					
OPT[dBm]	: -9.98	-9.98	-9.98	-15.09	NA	NO	0.00
	NA	NO					
OPR[dBm]	: -40.00	-40.00	-40.00	-30.00	NA	NO	8.00
	NA	NO					
CD[ps/nm]	: 0	0	0	-80000	-100	NO	100
	100	NO					
DGD[ps]	: 0.00	0.00	0.00	0.00	NA	NO	80.00
	NA	NO					
SOPMD[ps^2]	: 0.00	0.00	0.00	0.00	NA	NO	2000.00
	NA	NO					
OSNR[dB]	: 0.00	0.00	0.00	0.00	NA	NO	40.00
	NA	NO					
PDL[dB]	: 0.00	0.00	0.00	0.00	NA	NO	7.00
	NA	NO					
PCR[rad/s]	: 0.00	0.00	0.00	0.00	NA	NO	2500000.00
	NA	NO					
RX_SIG[dBm]	: -40.00	-40.00	-40.00	-30.00	NA	NO	1.00
	NA	NO					
FREQ_OFF[Mhz]	: 0	0	0	-3600	NA	NO	3600
	NA	NO					
SNR[dB]	: 0.00	0.00	0.00	7.00	NA	NO	100.00
	NA	NO					

```

Last clearing of "show controllers OPTICS" counters never
!

```

Performance Monitoring Configuration Example

This example shows how to enable performance monitoring and set PM thresholds and TCA reporting status on the coherent DSP controller:

```

Router#config
Router (config) #controller CoherentDSP0/2/0/16
Router (config-CoDSP) #perf-mon enable
Router (config-CoDSP) #pm 30-sec fec report Q max-tca enable
Router (config-CoDSP) #pm 30-sec fec report Q-margin max-tca enable
Router (config-CoDSP) #pm 30-sec fec report Q min-tca enable
Router (config-CoDSP) #pm 30-sec fec report Q-margin min-tca enable
Router (config-CoDSP) #pm 30-sec fec threshold Q max 1200
Router (config-CoDSP) #pm 30-sec fec threshold Q-margin max 500
Router (config-CoDSP) #pm 30-sec fec threshold Q min 900
Router (config-CoDSP) #pm 30-sec fec threshold Q-margin min 280
Router (config-CoDSP) #commit

```

Running Configuration

This example shows the running configuration on coherent DSP controllers:

```

Router#show run controller coherentdsp 0/2/0/16
Thu May 13 19:56:09.136 UTC
controller CoherentDSP0/2/0/16
  pm 30-sec fec report Q max-tca enable
  pm 30-sec fec report Q-margin max-tca enable
  pm 30-sec fec report Q min-tca enable
  pm 30-sec fec report Q-margin min-tca enable
  pm 30-sec fec threshold Q max 1200
  pm 30-sec fec threshold Q-margin max 500
  pm 30-sec fec threshold Q min 900
  pm 30-sec fec threshold Q-margin min 280
  perf-mon enable
!

```

Verification

This example shows how to verify the PM parameters on coherent DSP controllers. Verify the configuration changes in the highlighted fields:

```

Router#show controllers coherentdsp 0/2/0/16 pm current 30-sec fec
Thu May 27 23:04:54.167 UTC
g709 FEC in the current interval [23:04:30 - 23:04:54 Thu May 27 2021]
FEC current bucket type : Valid
  EC-BITS      : 0                               Threshold : 111484000000          TCA(enable) :
YES
  UC-WORDS    : 0                               Threshold : 5                       TCA(enable)  :
YES

```

Threshold	TCA	MIN	AVG	MAX	Threshold	TCA
(max)	(enable)				(min)	(enable)
PreFEC BER		:	0E-15	0E-15	0E-15	NO
0E-15	NO					
PostFEC BER		:	0E-15	0E-15	0E-15	NO
0E-15	NO					
Q[dB]		:	0.00	0.00	0.00	9.00 YES 120.00 YES
Q_Margin[dB]		:	0.00	0.00	0.00	2.80 YES 5.00 YES

```

!

```

Configuring Alarms Threshold

The alarms threshold can be configured for monitoring alarms on optics controllers:

Table 58: Alarms Threshold Parameters for Optics Controllers

Alarm Threshold Parameters	Description
CD	Sets the CD (chromatic dispersion) alarm threshold (cd-low-threshold and cd-high-threshold).
DGD	Sets the DGD (differential group delay) alarm threshold.
LBC	Sets the LBC (laser bias current) threshold in mA.
OSNR	Sets the OSNR (optical signal-to-noise ratio) alarm threshold.

Alarm Threshold Configuration Example

This example shows how to configure alarm threshold on the optics controller:

```
Router#config
Router(config)#controller optics 0/2/0/16
Router(config-Optics)#cd-low-threshold -2000
Router(config-Optics)#cd-high-threshold 2000
Router(config-Optics)#commit
```

Running Configuration

This example shows the running configuration on the optics controller:

```
Router#show run controller optics 0/2/0/16
Thu May 13 20:18:55.957 UTC
controller Optics0/2/0/16
  cd-low-threshold 2000
  cd-high-threshold 2000
!
```

Verification

This example shows how to verify the alarm threshold on optics controllers:

```
Router#show controller optics 0/2/0/16
Fri May 28 01:04:33.604 UTC
Controller State: Up
Transport Admin State: In Service
Laser State: Off
LED State: Off
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSFPDD 400G ZRP
  DWDM carrier Info: C BAND, MSA ITU Channel=61, Frequency=193.10THz,
  Wavelength=1552.524nm
  Alarm Status:
  -----
  Detected Alarms: None
  LOS/LOL/Fault Status:
  Alarm Statistics:
  -----
  HIGH-RX-PWR = 0          LOW-RX-PWR = 0
  HIGH-TX-PWR = 0          LOW-TX-PWR = 0
  HIGH-LBC = 0             HIGH-DGD = 0
  OOR-CD = 0               OSNR = 0
  WVL-OOL = 0              MEA = 0
  IMPROPER-REM = 0
```

```

TX-POWER-PROV-MISMATCH = 0
Laser Bias Current = 0.0 mA
Actual TX Power = -40.00 dBm
RX Power = -40.00 dBm
RX Signal Power = -40.00 dBm
Frequency Offset = 0 MHz
Laser Temperature = 0.00 Celsius
Laser Age = 0 %
DAC Rate = 1x1.25
Performance Monitoring: Enable
THRESHOLD VALUES

```

```

-----
Parameter                High Alarm  Low Alarm  High Warning  Low Warning
-----
Rx Power Threshold(dBm)   13.0       -24.0      10.0          -22.0
Tx Power Threshold(dBm)   0.0        -16.0      -2.0          -14.0
LBC Threshold(mA)         0.00       0.00       0.00          0.00
Temp. Threshold(celsius)  80.00      -5.00      75.00         0.00
Voltage Threshold(volt)   3.46       3.13       3.43          3.16

```

```
LBC High Threshold = 98 %
```

```
Configured Tx Power = -10.00 dBm
```

```
Configured CD High Threshold = -5000 ps/nm
```

```
Configured CD lower Threshold = -5000 ps/nm
```

```
Configured OSNR lower Threshold = 9.00 dB
```

```
Configured DGD Higher Threshold = 80.00 ps
```

```
Baud Rate = 60.1385459900 GBd
```

```
Modulation Type: 16QAM
```

```
Chromatic Dispersion 0 ps/nm
```

```
Configured CD-MIN -26000 ps/nm CD-MAX 26000 ps/nm
```

```
Second Order Polarization Mode Dispersion = 0.00 ps^2
```

```
Optical Signal to Noise Ratio = 0.00 dB
```

```
Polarization Dependent Loss = 0.00 dB
```

```
Polarization Change Rate = 0.00 rad/s
```

```
Differential Group Delay = 0.00 ps
```

```
Temperature = 21.00 Celsius
```

```
Voltage = 3.42 V
```

Transceiver Vendor Details

```

Form Factor                : QSFP-DD
Optics type                 : QSFPDD 400G ZRP
Name                       : CISCO-ACACIA
OUI Number                 : 7c.b2.5c
Part Number                : DP04QSDD-E30-19E
Rev Number                 : 10
Serial Number              : ACA244900GN
PID                        : QDD-400G-ZRP-S
VID                        : ES03
Firmware Version           : 161.06
Date Code(yy/mm/dd)       : 20/12/08

```

```
!
```

Configuring FEC Alarm Threshold

Table 59: Feature History Table

Feature Name	Release Information	Description
Configurable FDD and FED Alarm Threshold Values	Release 24.1.1	

Feature Name	Release Information	Description
		<p>We now ensure that you have accurate data to initiate proactive maintenance for non-critical FEC errors or take prompt action to prevent potential optical link data loss in your network. This is made possible because we've enabled the configuration of FEC (Forward Error Correction) Detected Degrade (FDD) alarm threshold values for non-critical FEC errors and FEC Excessive Degrade (FED) alarm threshold values for critical FEC errors. You can configure or clear these values for QDD-400G-ZR, QDD-400G-ZRP, and DP04QSDD-HE0 optical modules.</p> <p>Prior to this release, the router would automatically generate FEC alarms based on default threshold values.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <p>Modified the controller optics command by adding the following keywords:</p> <ul style="list-style-type: none"> • host fec-threshold excess-degrade raise • media fec-threshold excess-degrade raise • host fec-threshold excess-degrade clear • media fec-threshold excess-degrade clear • host fec-threshold detected-degrade raise • media fec-threshold detected-degrade raise • host fec-threshold detected-degrade clear • media fec-threshold detected-degrade clear

Feature Name	Release Information	Description
		<p>The fec-thresholds keyword is added to the show controller optics command.</p> <p>YANG Data Model:</p> <ul style="list-style-type: none"> • New XPaths for Cisco-IOS-XR-controller-optics-oper.yang • Cisco-IOS-XR-controller-optics-fec-thresholds.yang

Forward Error Correction (FEC) is used to control errors during data transmission. FEC works by adding data redundancy to the transmitted message. This redundancy allows the receiver to detect and correct a limited number of errors occurring anywhere in the message, instead of the transmitter resending the entire message. For additional information on FEC, see [Understanding FEC and Its Implementation](#).

There are two types of FEC alarms:

- **FEC Detected Degrade (FDD) alarm:** The FDD alarm is raised when the link degradation is within the permissible limit and does not cause traffic disruption. This alarm indicates the system is working harder than usual to maintain data transmission. Link degradation could be due to issues in the cable, network congestion, or other hardware failure.
- **FEC Excessive Degrade (FED) alarm:** The FED alarm is raised when the link degradation exceeds beyond the permissible limit and causes traffic disruption. This alarm indicates the system is working harder than usual to maintain data transmission. Without corrective measures, network performance deteriorates further and eventually results in traffic loss. Link degradation could be due to issues in the cable, network congestion, or other hardware failure.

The FEC alarms threshold values can now be configured to control alarms (raise and clear FEC alarms) on both media and host side of the optical transceiver. The optical transceiver is divided into two sides, the host side, which is positioned towards the router, and the media side, which is positioned towards the wire or cable media.

When the average bit error rate (BER) exceeds the **raise threshold value**, the FEC alarm is raised (or asserted). Similarly, when the BER drops below the **clear threshold value**, then the alarm is cleared (or de-asserted).

Guidelines and Restrictions for Setting the FEC Alarm Thresholds

- The **raise threshold value** must always be greater than the **clear threshold value** for both FDD and FED alarms.
- The **raise or clear threshold value** of FED alarm must always be greater than the **raise or clear threshold value** of the FDD alarm.
- While the router configuration permits a range of 1 to 18446744073709551615, the router only supports a range of 1 to 204600000000000000. The threshold value provided by users is converted from a 64 bit number to a 16 bit number. As a result, there is minor variation between the user provided value (configured value) and the programmed value. The user input (threshold value) is appended with exponents relative to E-18.

Table 60:

Configured Value	Programmed Value (Displayed using the Show CLI command)	Pattern	
1, 2, 3,,10	0, 1, 2,.....,9	1<ConfiguredValue< 10, show command value = ConfiguredValue - 1	1->>>0.9999, displayed as 0 and so on
11,12,13,.....,99	1.0, 1.1, 1.2,...9.8	10<ConfiguredValue<99, show command value = ConfiguredValue - 0.1	
111,222,333...999	1.10, 2.21, 3.32	100<ConfiguredValue<999, show command value = ConfiguredValue - 0.01	
1111,1112,1113 upto 2047	1.110, 1.111, 1.112...	1000<ConfiguredValue< 2047 show command value = ConfiguredValue - 0.001	
2050, 12345, 23456,65432,...	2.0500, 1.2300, 2.3400,6.5400...	2047<ConfiguredValue<maximum- range show command value = first 3 digits appended by 0s	

Configuration Examples to Set FEC Alarm Threshold

Examples to configure FEC alarm threshold:

- [Configuring FED Alarm Thresholds, on page 301](#)
- [Configuring FDD Alarm Thresholds, on page 300](#)

Configuring FDD Alarm Thresholds

FDD Configuration Example

This example shows how to set FDD clear and raise alarm thresholds on the optics controller:

```
Router#config
Router(config)#controller optics 0/0/0/10
Router(config-Optics)#host fec-threshold detected-degrade clear 12000
Router(config-Optics)#host fec-threshold detected-degrade raise 22000
Router(config-Optics)#commit
Router(config-Optics)#end
```

Running Configuration

This example shows the running configuration on the optics controller:

```
Router#show running-config controller optics 0/0/0/10
Sat Feb 3 06:01:56.354 UTC
controller Optics0/0/0/10
host fec-threshold detected-degrade raise 22000
host fec-threshold detected-degrade clear 12000
!
!
```

Verification

This example shows how to verify the alarm threshold values on optics controllers:

```
Router#show controller optics 0/0/0/10 fec-thresholds
FEC Threshold Information
```

	Raise	Clear
Media FEC excess degrade	1.2600E-02	1.2100E-02
Media FEC detected degrade	1.1700E-02	1.1300E-02
Host FEC excess degrade	2.4000E-02	2.4000E-03
Host FEC detected degrade	2.2000E-14	1.1989E-14

Configuring FED Alarm Thresholds

FED Configuration Example

This example shows how to set FED raise and clear alarm thresholds on the optics controller:

```
Router#config
Router(config)#controller optics 0/0/0/12
Router(config-Optics)#host fec-threshold excess-degrade clear 14000
Router(config-Optics)#host fec-threshold excess-degrade raise 24000
Router(config-Optics)#commit
Router(config-Optics)#end
```

Running Configuration

This example shows the running configuration on the optics controller:

```
Router#show running-config controller optics 0/0/0/12
Sat Feb 3 06:02:00.153 UTC
controller Optics0/0/0/12
host fec-threshold excess-degrade raise 24000
host fec-threshold excess-degrade clear 14000
!
```

Verification

This example shows how to verify the alarm threshold values on optics controllers:

```
Router#show controller optics 0/0/0/12 fec-thresholds
FEC Threshold Information
```

	Raise	Clear
Media FEC excess degrade	1.2600E-02	1.2100E-02
Media FEC detected degrade	1.1700E-02	1.1300E-02
Host FEC excess degrade	2.3900E-14	1.3999E-14
Host FEC detected degrade	9.0000E-03	9.0000E-04

Media Link-down PreFEC Degradation Enablement

Table 61: Feature History Table

Feature Name	Release Information	Description
Media Link-down PreFEC Degradation Enablement	Release 24.3.1	<p>Introduced in this release on: Fixed Systems(8200, 8700); Centralized Systems (8600); Modular Systems (8800 [LC ASIC: Q100, Q200, P100]).</p> <p>The Media Link-down PreFEC Degradation functionality can be used to protect the media side of the optical transceiver during transmission errors.</p> <p>By using this feature, you can proactively switch the traffic to standby path when the BER counter crosses the threshold value. This feature helps to avoid further traffic impact when the optical network reaches more noise or error.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <p>Modified the <code>controller optics</code> command by adding the media link-down prefec-degrade keyword.</p> <p>YANG Data Model:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-controller-optics-oper.yang</code> • New XPath for <code>Cisco-IOS-XR-um-cont-optics-fec-threshold-cfg.yang</code> <p>(see GitHub, YANG Data Models Navigator)</p>

The Media Link-down PreFEC Degradation functionality can be used to protect the media side of the optical transceiver during transmission errors, such as errors due to noise, or data transmission errors. This feature is disabled by default. You can enable this feature by using the **media link-down prefec-degrade** command.

Prerequisites for using Media Link-down PreFEC Degradation Functionality

To use the Media Link-down PreFEC Degradation functionality, you must configure the FEC Alarm Threshold. For information on configuring FEC alarms threshold, see [Configuring FEC Alarm Threshold](#).

About Media Link-down PreFEC Degradation Functionality

Prior to this release, the FEC Alarm Threshold functionality enabled you to configure the FEC alarms threshold values to control alarms (raise and clear FEC alarms) on media and host side of the optical transceiver. Using the FEC Alarm Threshold functionality, you can configure the FDD and FED alarm threshold values and set the **raise threshold value** and **clear threshold value** values to control alarms.

After you configure FEC Alarm Threshold and enable Media Link-down PreFEC Degradation functionality, you get the alarm notification when the average bit error rate (BER) exceeds the threshold value. This triggers

link-down and enables switchover functionality automatically. The traffic is switched to standby path, and remains in the standby path until the alarm is cleared or based on the settings done by the network operator.



Note In Cisco IOS XR Release 24.3.1, the Link-down PreFEC Degrade feature is supported only on the media side of the optical transceiver.

Configure Media Link-down PreFEC Degrade

The purpose of this task is to enable the media link-down preFEC degrade functionality to proactively switch the traffic to standby path.

Procedure

Step 1 Execute the **media link-down prefec-degrade** command to configure link-down preFEC degrade on the media side of the optics controller.

Example:

```
Router#config
Router(config)#controller optics 0/2/0/16
Router(config-Optics)#media link-down prefec-degrade
Router(config-Optics)#commit
```

Step 2 Execute the **show running-config controller optics R/S/I/P** command to view the running configuration on the optics controller.

Example:

```
Router#show running-config controller optics 0/2/0/16
Thu May 13 20:18:55.957 UTC
controller Optics0/2/0/16
  media link-down prefec-degrade
!
```

Step 3 Execute the **show controller optics R/S/I/P** command to verify link-down preFEC degrade feature on optics controllers.

Example:

```
Router#show controller optics 0/2/0/16
Fri May 28 01:04:33.604 UTC
Controller State: Up
Transport Admin State: In Service
Laser State: On
Media linkdown prefec degrade : Enabled
LED State: Green
FEC State: FEC ENABLED
Optics Status
  Optics Type: QSPDD 400G ZRP
  DWDM carrier Info: C BAND, MSA ITU Channel=61, Frequency=193.10THz,
  Wavelength=1552.524nm
  Alarm Status:
  -----
  Detected Alarms: None
  LOS/LOL/Fault Status:
  Alarm Statistics:
  -----
```

```

HIGH-RX-PWR = 0          LOW-RX-PWR = 0
HIGH-TX-PWR = 0          LOW-TX-PWR = 0
HIGH-LBC = 0            HIGH-DGD = 0
OOR-CD = 0              OSNR = 0
WVL-OOL = 0             MEA = 0

```

```

IMPROPER-REM = 0
TX-POWER-PROV-MISMATCH = 0
Laser Bias Current = 0.0 mA
Actual TX Power = -40.00 dBm
RX Power = -40.00 dBm
RX Signal Power = -40.00 dBm
Frequency Offset = 0 MHz
Laser Temperature = 0.00 Celsius
Laser Age = 0 %
DAC Rate = 1x1.25
Performance Monitoring: Enable

```

THRESHOLD VALUES

```

-----
Parameter                High Alarm  Low Alarm  High Warning  Low Warning
-----
Rx Power Threshold(dBm)   13.0       -24.0      10.0          -22.0
Tx Power Threshold(dBm)   0.0        -16.0      -2.0          -14.0
LBC Threshold(mA)         0.00       0.00      0.00          0.00
Temp. Threshold(celsius)  80.00      -5.00     75.00         0.00
Voltage Threshold(volt)   3.46       3.13      3.43          3.16

```

```

LBC High Threshold = 98 %
Configured Tx Power = -10.00 dBm
Configured CD High Threshold = -5000 ps/nm
Configured CD lower Threshold = -5000 ps/nm
Configured OSNR lower Threshold = 9.00 dB
Configured DGD Higher Threshold = 80.00 ps
Baud Rate = 60.1385459900 GBd
Modulation Type: 16QAM
Chromatic Dispersion 0 ps/nm
Configured CD-MIN -26000 ps/nm CD-MAX 26000 ps/nm
Second Order Polarization Mode Dispersion = 0.00 ps^2
Optical Signal to Noise Ratio = 0.00 dB
Polarization Dependent Loss = 0.00 dB
Polarization Change Rate = 0.00 rad/s
Differential Group Delay = 0.00 ps
Temperature = 21.00 Celsius
Voltage = 3.42 V

```

Transceiver Vendor Details

```

Form Factor           : QSFP-DD
Optics type           : QSFPDD 400G ZRP
Name                  : CISCO-ACACIA
OUI Number            : 7c.b2.5c
Part Number           : DP04QSDD-E30-19E
Rev Number            : 10
Serial Number         : ACA244900GN
PID                   : QDD-400G-ZRP-S
VID                   : ES03
Firmware Version      : 161.06
Date Code(yy/mm/dd)  : 20/12/08

```

!

Alarms Troubleshooting

Table 62: Feature History Table

Feature Name	Release	Description
Enhanced Alarm Prioritization, Monitoring, and Management	Release 24.1.1	To improve alarm prioritization, precise monitoring, and streamlined management, ensuring a more efficient and responsive approach to network events, we have introduced the suppression of LOL (Loss of Line) alarms when the LOS-P (Loss of Signal-Payload) alarm is generated, the clearing of alarm static counters, and the suppression of warnings when the respective alarm is triggered.

In this latest release, we bring forth advanced features to elevate your alarm management experience. The key highlights include:

- **LOL Alarm Suppression:** Prioritize the detection and handling of critical LOS-P (Loss of Signal-Payload) alarms by suppressing LOL (Loss of Line) alarms when LOS-P alarms are generated.
- **Clearing Static Counters:** Gain the ability to clear alarm static counters using the command `clear counters controller coherentDSP location`, facilitating focused monitoring within a specific time frame.
- **Warning Suppression:** Enhance your monitoring environment by suppressing warnings when the corresponding alarm is triggered, effectively preventing redundant or repetitive alerts.

These enhancements collectively contribute to improved alarm prioritization, precise monitoring, and streamlined management, ensuring a more efficient and responsive approach to network events.

This section contains the procedures for troubleshooting alarms.

CD Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The Chromatic Dispersion (CD) alarm is raised when the detected chromatic dispersion value is above or below the configured threshold values.

Clear the CD Alarm

Procedure

Configure threshold value within range if CD value is not within the threshold range.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

DGD Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The Differential Group Delay (DGD) alarm is raised when the value of the differential group delay read by the pluggable port module exceeds the configured threshold value.

Clear the DGD Alarm

Procedure

Configure the threshold value within range if DGD value is not within the threshold range.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

FLEXO_LOF

Default Severity: Critical

Logical Object: OTN

Flexo LOF alarm is raised when loss of alignment is detected on the Flexo frame for more than 3ms.

Clear the FLEXO_LOF Alarm

Procedure

Step 1 Identify and correct the underlying cause of mis-alignment. The Flexo LOF (Loss of Frame) alarm is cleared

Step 2 when good alignment is detected on the Flexo frame for more than 3ms.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

FLEXO_LOM

Default Severity: Critical

Logical Object: OTN

Flexo LOM (Loss of Multi-Frame) is raised when loss of multi-frame alignment is detected on the Flexo multi-frame for more than 10ms

Clear the FLEXO_LOM Alarm

Procedure

Step 1 Identify and correct the underlying cause of mis-alignment. The Flexo LOM alarm is cleared when good

Step 2 multi-frame alignment is detected on the Flexo multi-frame.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-LASERBIAS Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The HI-LASERBIAS alarm is raised when the physical pluggable port laser detects a laser bias value beyond the configured high threshold.

Clear the HI-LASERBIAS Alarm

Procedure

Configure the threshold value within range if high laser bias threshold value is not within the threshold range.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-RXPOWER Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The HI-RXPOWER alarm occurs on the client optics controller when the measured individual lane optical signal power of the received signal exceeds the default threshold. The HI-RXPOWER alarm occurs on the trunk optics controller when the total optical signal power of the received signal exceeds the default threshold.

Clear the HI-RXPOWER Alarm

Procedure

Physically verify by using a standard power meter that the optical input power is overcoming the expected power threshold. Connect an attenuator accordingly.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-RXPOWER Warn

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Software

The HI-RXPOWER warning occurs on the client optics controller when the measured individual lane optical signal power of the received signal exceeds the default threshold. The HI-RXPOWER warning occurs on the trunk optics controller when the total optical signal power of the received signal exceeds the default threshold.

Clear the HI-RXPOWER Warn Alarm

Procedure

Physically verify by using a standard power meter that the optical input power is overcoming the expected power threshold. Connect an attenuator accordingly.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-TEMP Alarm

Default Severity: Critical

Logical Object: Software

The HI-TEMP alarm occurs when the optical module temperature exceeds the default threshold.

Clear the HI-TEMP Alarm

Procedure

Verify the fan is intact and empty slots are blocked for cooling

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-TEMP Warn

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Software

The HI-TEMP warning occurs when the optical module temperature exceeds the default threshold.

Clear the HI-TEMP Warn Alarm

Procedure

Verify the fan is intact and empty slots are blocked for cooling

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-TXPOWER Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The HI-TXPOWER alarm occurs on the client optics controller when the measured individual lane optical signal power of the transmitted signal exceeds the default threshold. The HI-TXPOWER alarm occurs on the trunk optics controller when the total optical signal power of the transmitted signal exceeds the default threshold.

Clear the HI-TXPOWER Alarm

Procedure

Physically verify by using a standard power meter that the optical output power is overcoming the expected power threshold.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

HI-TXPOWER Warn

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Software

The HI-TXPOWER warning occurs on the client optics controller when the measured individual lane optical signal power of the transmitted signal exceeds the default threshold. The HI-TXPOWER warning occurs on the trunk optics controller when the total optical signal power of the transmitted signal exceeds the default threshold.

Clear the HI-TXPOWER Warn Alarm

Procedure

Physically verify by using a standard power meter that the optical output power is overcoming the expected power threshold.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

IMPROPER-REM

Default Severity: Critical

Logical Object: Software

The Improper Removal alarm is raised when a physical pluggable is not present on a service-provisioned port.

Clear the IMPROPER-REM Alarm

Procedure

Insert the appropriate QSFP.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LOF

Default Severity: Critical

Logical Object: OTN

Flexo LOF alarm is raised when loss of alignment is detected on the Flexo frame for more than 3ms.

Clear the LOF Alarm

Procedure

Step 1 Identify and correct the underlying cause of mis-alignment. The Flexo LOF (Loss of Frame) alarm is cleared

Step 2 when good alignment is detected on the Flexo frame for more than 3ms.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LOL

Default Severity: Major

Logical Object: Software

Loss of lock (LOL) alarm is raised when Loss of lock is detected on the receive side of the CDR (Clock and Data Recovery)

Clear the LOL Alarm

Procedure

Verify the fiber and power levels.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LOM

Default Severity: Critical

Logical Object: OTN

Flexo LOM (Loss of Multi-Frame) is raised when loss of multi-frame alignment is detected on the Flexo multi-frame for more than 10ms

Clear the LOM Alarm

Procedure

Step 1 Identify and correct the underlying cause of mis-alignment. The Flexo LOM alarm is cleared when good

Step 2 multi-frame alignment is detected on the Flexo multi-frame.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LO-RXPOWER Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The LO-RXPOWER alarm is raised on the client or trunk optics controller when the measured individual lane optical signal power of the received signal falls below the default threshold.

Clear the LO-RXPOWER Alarm

Procedure

Verify that the trunk-rx port is cabled correctly, and clean the fiber connecting the faulty TXP/MXP card to the drop port of the DWDM card.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LO-RXPOWER Warn

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Software

The LO-RXPOWER warning is raised on the client or trunk optics controller when the measured individual lane optical signal power of the received signal falls below the default threshold.

Clear the LO-RXPOWER Warn Alarm

Procedure

Verify that the trunk-rx port is cabled correctly, and clean the fiber connecting the faulty TXP/MXP card to the drop port of the DWDM card.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LOS

Default Severity: Major

Logical Object: Software

This alarm occurs when there is a loss of signal

Clear the LOS Alarm

Procedure

Identify and correct the underlying cause of signal LOS. The alarm is cleared when signal is improved.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LOS-P

Default Severity: Minor

Logical Object: OTN

This alarm occurs when there is a loss of signal.

Clear the LOS-P Alarm

Procedure

Identify and correct the underlying cause of signal LOS. The alarm is cleared when signal is improved.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LO-TXPOWER Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The LO-TXPOWER alarm is raised on the client or trunk optics controller when the measured individual lane optical signal power of the transmitted signal falls below the default threshold.

Clear the LO-TXPOWER Alarm

Procedure

Verify the optics detection and any failures.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

LO-TXPOWER Warn

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Software

The LO-TXPOWER warning is raised on the client or trunk optics controller when the measured individual lane optical signal power of the transmitted signal falls below the default threshold.

Clear the LO-TXPOWER Warn Alarm

Procedure

Verify the optics detection and any failures.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

OOR_CD

Default Severity: Minor

Logical Object: Controller

This alarm occurs when the Chromatic Dispersion is out of range

Clear the OOR_CD Alarm

Procedure

Configure threshold value within range if CD value is not within the threshold range.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

OSNR Alarm

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The Optical Signal Noise Ratio (OSNR) alarm occurs when the measured OSNR falls below the threshold.

Clear the OSNR Alarm

Procedure

- Step 1** Verify the value of the minimum acceptable OSNR value of Cisco 8000 using the show controller optics R/S/I/P command.
 - Step 2** If the value is not within the OSNR threshold range, configure the minimum acceptable OSNR value using the controller optics R/S/I/P osnr-low-threshold command in the config mode. The range is 0 to 4000 (in units of 01db).
 - Step 3** If the value is within the range of the minimum acceptable OSNR, contact TAC .
If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).
-

UNC-WORD Alarm

Default Severity: Not Reported (NR), Not-Alarmed, Non-Service-Affecting (NSA)

Logical Object: OTN

The Uncorrected FEC Word (UNC-WORD) condition is raised when the FEC is unable to correct the frame.

Clear the UNC-WORD Alarm

Procedure

- Step 1** Ensure that the fiber connector for the card is completely plugged in.
- Step 2** Ensure that the ports on the far end and near end nodes have the same port rates and FEC settings.
- Step 3** If the BER threshold is correct and at the expected level, use an optical test set to measure the power level of the line to ensure it is within guidelines. For specific procedures to use the test set equipment, consult the manufacturer.
- Step 4** If the optical power level is good, verify that the optical receive levels are within the acceptable range.
- Step 5** If receive levels are good, clean the fibers at both ends.
- Step 6** If the condition does not clear, verify that a single-mode fiber is used.

Step 7 Verify if the fiber is of single-mode type.

Step 8 Clean the fiber connectors at both ends for a signal degrade.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

WVL-00L

Default Severity: Major

Logical Object: Controller

The Wavelength Out of Lock alarm is raised when the port detects the optical input frequency to be out of range.

Clear the WVL-00L Alarm

Procedure

Step 1 1 Verify the wavelength configuration.

Step 2 2 Verify whether the pluggable is inserted properly.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).



CHAPTER 16

Configuring Controllers

This chapter describes the Optics Controller for the 36-port QSFP56-DD 400 GbE and 48-port QSFP28 100 GbE Line Cards. This chapter also describes the procedures used to configure the controllers.



Note When two MACsec enabled Cisco 8000 Series Routers with Coherent Line Cards are connected, there is no compatibility between Coherent Line Cards of IOS XR Release.

- breakout - Configure breakout mode ('breakout 4x10' only.)
- clear - Clear the uncommitted configuration.
- commit - Commit the configuration changes to running.
- do - Run an exec command.
- end - Exit from configure mode.
- exit - Exit from this submode.
- ext-description - Set ext-description for this controller.
- no - Negate a command or set its defaults.
- pwd - Commands used to reach current submode.
- root - Exit to the global configuration mode.
- show - Show contents of configuration.

Following controller configuration options are supported on the router:

- [How to Configure Controllers, on page 318](#)
- [Diagnostic Parameters for Optical Transceivers, on page 321](#)
- [View Coherent Optical Transceiver Module Parameters , on page 326](#)
- [Display of Alarms for Coherent Optical Transceivers, on page 329](#)
- [Restrictions, on page 330](#)
- [Verification, on page 330](#)
- [Loopback on Optical Transceivers, on page 332](#)
- [Media Side Input Loopback Configuration, on page 334](#)
- [Media Side Output Loopback, on page 335](#)

- [Host Side Input Loopback Configuration, on page 336](#)
- [Host Side Output Loopback Configuration, on page 338](#)

How to Configure Controllers

This section contains the following procedures:

Configuring Optics Controller

Configuring optics controller of breakout 4x10:

```
RP/0/RP0/CPU0:uut#configure
Fri Oct 11 16:22:31.222 UTC
RP/0/RP0/CPU0:uut(config)#controller optics 0/1/0/28
RP/0/RP0/CPU0:uut(config-Optics)#breakout 4x10
RP/0/RP0/CPU0:uut(config-Optics)#commit
Fri Oct 11 16:23:26.868 UTC
RP/0/RP0/CPU0:uut(config-Optics)#end
RP/0/RP0/CPU0:uut#
RP/0/RP0/CPU0:uut#show running-config controller optics 0/1/0/28
Fri Oct 11 16:23:41.273 UTC
controller Optics0/1/0/28
breakout 4x10
!
```

Disabling Optical Modules

This feature provides the ability to disable and re-enable an optical module through CLI, which simulates online insertion and removal (OIR) by disabling power to the transceiver port.

Typical troubleshooting procedures for optical modules can include performing OIR by removing and re-installing the module, which requires onsite personnel to physically reseal the optical module. The ability to remotely disable and enable an optical module can significantly reduce operational expenses.

Example

The following output shows a QSFP28 module powered on and in UP state:

```
Router# show controllers optics 0/0/0/0

Controller State: Up

Transport Admin State: In Service

Laser State: Off

LED State: Not Applicable

FEC State: FEC ENABLED

Optics Status

    Optics Type:  QSFP28 100G FR
    Wavelength = 1311.00 nm

    Alarm Status:
    -----
```


Detected Alarms: None

LOS/LOL/Fault Status:

Laser Bias Current = 26.2 mA
Actual TX Power = 0.73 dBm
RX Power = -0.68 dBm

Performance Monitoring: Disable

THRESHOLD VALUES

Parameter	High Alarm	Low Alarm	High Warning	Low Warning
Rx Power Threshold(dBm)	7.4	-10.4	4.5	-6.3
Tx Power Threshold(dBm)	7.0	-6.3	4.0	-2.4
LBC Threshold(mA)	100.00	8.00	83.00	10.00
Temp. Threshold(celsius)	75.00	-5.00	70.00	0.00
Voltage Threshold(volt)	3.63	2.97	3.46	3.13

Polarization parameters not supported by optics

Temperature = 27.92 Celsius
Voltage = 3.24 V

Transceiver Vendor Details

```

Form Factor           : QSFP28
Optics type          : QSFP28 100G FR
Name                  : CISCO-CISCO
OUI Number           : 00.00.0c
Part Number           : 10-3248-01
Rev Number            : 01
Serial Number         : FBN2331A114
PID                   : QSFP-100G-FR-S
VID                   : ESO
Date Code(yy/mm/dd)  : 19/09/19
  
```

To disable the module, use the **transceiver disable** command in controller optics configuration mode:

```

Router(config)# controller optics 0/0/0/0
Router(config-Optics)# transceiver disable
Router(config-Optics)# commit
Router(config-Optics)# end
  
```

The following example shows the QSFP28 module disabled and powered down:

```
Router# show controllers optics 0/0/0/0
```

Controller State: **Down**

Transport Admin State: In Service

Laser State: Off

Optics Status

Optics Type: **Unknown optics**
Wavelength = 0.00 nm

Alarm Status:

```

-----
Detected Alarms: None

LOS/LOL/Fault Status:

TX Power = N/A
RX Power = N/A

Performance Monitoring: Disable

THRESHOLD VALUES
-----

Parameter                High Alarm  Low Alarm  High Warning  Low Warning
-----
Rx Power Threshold(dBm)    7.4        -10.4     4.5           -6.3
Tx Power Threshold(dBm)    7.0        -6.3     4.0           -2.4
LBC Threshold(mA)         100.00     8.00     83.00        10.00
Temp. Threshold(celsius)  75.00     -5.00    70.00         0.00
Voltage Threshold(volt)    3.63       2.97     3.46         3.13

Polarization parameters not supported by optics

Temperature = 0.00 Celsius
Voltage = 0.00 V

Transceiver Vendor Details

```

To re-enable the module, use the **no transceiver disable** command in controller optics configuration mode.

Diagnostic Parameters for Optical Transceivers

Table 63: Feature History Table

Feature Name	Release Information	Description
Diagnostic Parameters for Optical Transceivers	Release 7.11.1	

Feature Name	Release Information	Description
		<p>You can analyze the diagnostic parameters for optical transceivers installed on a network device and detect potential issues with the optical transceivers, such as excessive power levels, abnormal temperature readings, or degradation of the optical signal. Such analysis is possible because the show controllers optics command now displays the following diagnostic parameters:</p> <ul style="list-style-type: none"> • Effective Signal to Noise Ratio (eSNR) • Pulse Amplitude Modulation with Four Levels (PAM4) Level Transition Parameter (LTP) • Pre-Forward Error Correction (FEC) and Post-FEC Bit Error Rate (BER) • Frame Error Count (FERC) • Laser age • Thermoelectric Cooler (TEC) current • Laser frequency • Laser temperature <p>For additional information on VDM (Versatile Diagnostics Monitoring), see the Common Management Interface Specification.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <ul style="list-style-type: none"> • The observable-info keyword is added to the show controller optics command. <p>YANG Data Model:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-controller-optics-oper.yang</code>

Feature Name	Release Information	Description
		(see GitHub , YANG Data Models Navigator)

In order to monitor and report the performance of an optical transceiver and thereby enhancing the troubleshooting capabilities of the optical transceiver, the **observable-info** keyword is added to the **show controllers optics** command to display the diagnostics parameters. These parameters help in monitoring the health of the network when the optical transceiver heats up, when the link is down, when alarms are raised, or when there's traffic loss in the network. This improvement in the **show controllers optics** command now displays the following diagnostic parameters:

- Effective Signal to Noise Ratio (eSNR)
- Pulse Amplitude Modulation with Four Levels (PAM4) Level Transition Parameter (LTP)
- Pre-Forward Error Correction (FEC) and Post-FEC Bit Error Rate (BER)
- Frame Error Count (FERC)
- Laser age
- Thermoelectric Cooler (TEC) current
- Laser frequency
- Laser temperature



Note Not all optical transceivers support the **observable-info** keyword. Also, the parameters that are displayed depend on what the optical transceiver supports, that is, not all optical transceivers display the same parameters. For additional information on VDM (Versatile Diagnostics Monitoring), see the [Common Management Interface Specification](#).

Verification

The following **show controllers optics observable-info** command displays the monitoring parameters of the optical transceiver present in the 0/0/0/9 location ID. The 0/0/0/9 location ID represents rack/slot/instance/port. Based on the requirement, the network administrators can use the displayed values of this command for monitoring and troubleshooting.

```
Router#show controllers optics 0/0/0/9 observable-info
Observable Information
```

```
[eSNR Media Input]
Unit: dB
Id      Value      TCAWarn      LowThreshWarn      HighThresWarn      LowThreshAlarm
HighThreshAlarm      Low High      Low High
Lane0   21.30      n n          0.00              0.00              0.00
0.00
Lane1   22.05      n n          0.00              0.00              0.00
0.00
Lane2   22.62      n n          0.00              0.00              0.00
0.00
Lane3   22.05      n n          0.00              0.00              0.00
```

```

0.00          n  n    n  n

[PAM4 Level Transition Parameter Media Input]
Unit: dB
Id      Value          LowThreshWarn    HighThresWarn    LowThreshAlarm
HighThreshAlarm    TCAWarn    TCAAlarm

          Low High    Low High
Lane0    47.79          0.00          0.00          0.00
0.00          n  n    n  n
Lane1    54.70          0.00          0.00          0.00
0.00          n  n    n  n
Lane2    64.34          0.00          0.00          0.00
0.00          n  n    n  n
Lane3    59.64          0.00          0.00          0.00
0.00          n  n    n  n

[Pre-FEC BER Minimum Media Input]
Unit: n/a
Id      Value          LowThreshWarn    HighThresWarn    LowThreshAlarm
HighThreshAlarm    TCAWarn    TCAAlarm

          Low High    Low High
Module  0.000E+00      0.000E+00      0.000E+00      0.000E+00
0.000E+00      n  n    n  n

[Pre-FEC BER Minimum Host Input]
Unit: n/a
Id      Value          LowThreshWarn    HighThresWarn    LowThreshAlarm
HighThreshAlarm    TCAWarn    TCAAlarm

          Low High    Low High
Module  0.000E+00      0.000E+00      0.000E+00      0.000E+00
0.000E+00      n  n    n  n

[Pre-FEC BER Maximum Media Input]
Unit: n/a
Id      Value          LowThreshWarn    HighThresWarn    LowThreshAlarm
HighThreshAlarm    TCAWarn    TCAAlarm

          Low High    Low High
Module  0.000E+00      0.000E+00      0.000E+00      0.000E+00
0.000E+00      n  n    n  n

[Pre-FEC BER Maximum Host Input]
Unit: n/a
Id      Value          LowThreshWarn    HighThresWarn    LowThreshAlarm
HighThreshAlarm    TCAWarn    TCAAlarm

          Low High    Low High
Module  0.000E+00      0.000E+00      0.000E+00      0.000E+00
0.000E+00      n  n    n  n

[Pre-FEC BER Average Media Input]
Unit: n/a
Id      Value          LowThreshWarn    HighThresWarn    LowThreshAlarm
HighThreshAlarm    TCAWarn    TCAAlarm

          Low High    Low High
Module  0.000E+00      0.000E+00      0.000E+00      0.000E+00
0.000E+00      n  n    n  n

[Pre-FEC BER Average Host Input]
Unit: n/a

```

Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[Pre-FEC BER Current Media Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[Pre-FEC BER Current Host Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[FERC Minimum Media Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[FERC Minimum Host Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[FERC Maximum Media Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[FERC Maximum Host Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm
HighThreshAlarm		TCAWarn	TCAAlarm		
		Low High	Low High		
Module	0.000E+00		0.000E+00	0.000E+00	0.000E+00
0.000E+00		n n	n n		
[FERC Average Media Input]					
Unit: n/a					
Id	Value		LowThreshWarn	HighThresWarn	LowThreshAlarm

```

HighThreshAlarm      TCAWarn  TCAAlarm

                    Low High  Low High
Module  0.000E+00    n  n    0.000E+00    0.000E+00
0.000E+00

[FERC Average Host Input]
Unit: n/a
Id      Value
HighThreshAlarm  TCAWarn  TCAAlarm

                    Low High  Low High
Module  0.000E+00    n  n    0.000E+00    0.000E+00
0.000E+00

[FERC Current Media Input]
Unit: n/a
Id      Value
HighThreshAlarm  TCAWarn  TCAAlarm

                    Low High  Low High
Module  0.000E+00    n  n    0.000E+00    0.000E+00
0.000E+00

[FERC Current Host Input]
Unit: n/a
Id      Value
HighThreshAlarm  TCAWarn  TCAAlarm

                    Low High  Low High
Module  0.000E+00    n  n    0.000E+00    0.000E+00
0.000E+00

```

View Coherent Optical Transceiver Module Parameters

Table 64: Feature History Table

Feature Name	Release Information	Description
View Coherent Optical Transceiver Module Parameters	Release 24.1.1	<p>You can now monitor the overall functioning and status of the coherent optical transceiver. This is possible because you can now view the module state and data path state of the optical transceivers, which give you an insight into the current state of the optical transceiver.</p> <p>This feature modifies the output of the show controller optics command.</p>

The display of the [show controller optics](#) command is now improved to include information about the module state and datapath state of the optical transceivers. You can view these details using the **information all** and **information counters** keywords of the [show controller optics](#) command.

Module State Machine (MSM) defines host-module interactions and behavioral characteristics of the optical module, such as the initialization of the management interface and the module power mode. The Module State field provides the current status of the optical transceiver. The optical transceiver can be in any one of the following states:

- Low power
- Power up
- Ready
- Power down
- Fault

Data Path State Machine (DPSM) defines the host-module interactions and behavioral characteristics needed for the initialization of one particular data path, which represents the signal flow and signal processing of any one instance of one type of application.

The Datapath State field provides the current state of the data path on each host lane. It represents the initialization status of the resources associated with a data path in response to host configuration settings or commands.

By default, the data paths begin initializing when the module state is ready. The data path state is dependent on the module state. The host uses the activated data path to carry traffic.



Note For additional information on the module state and datapath state, see the [Common Management Interface Specification](#).

Verification

The **show controllers optics r/s/i/p information counters** command displays the parameters of the optical transceiver present in the 0/0/0/8 location ID. The 0/0/0/8 location ID represents rack/slot/instance/port. Based on the requirement, the network administrators can use the displayed values of this command for monitoring and troubleshooting. This example displays the details when the optical transceiver operates in the transponder mode.

```
Router#show controllers optics 0/0/0/8 information counters
Fri Feb 16 11:06:31.415 UTC

Module State : Ready

Datapath State [Client-0]: TX Turn On

Acquisition Counter:      INVALID

HOST SIDE ALARM COUNTERS
=====
Host-Intf-0-FDD-Alarm-Counter[0]                Host-Intf-0-FED-Alarm-Counter[0]

HOST SIDE FEC-BER FEC-FERC CURRENT VALUES
=====
Host-Intf-0-FEC-BER[0.00E+00]                    Host-Intf-0-FEC-FERC[0.00E+00]

Supported Loopback Types :
```

```

=====
[1.] Media Internal
[2.] Media Line
[3.] Host Line
[4.] Host Internal
[5.] Host Per Lane
[6.] Media Per Lane
[7.] Simultaneous Media Host

Unsupported Loopback Types :
=====

Media Configured Loopback : Media Loopback None
Media Applied Loopback    : Media Loopback None

Host Configured Loopback  : Host Loopback None
Host Applied Loopback     : Host Loopback None

```

```

FW Upgrade Capability Mode:
=====
Supports Both Warm & Cold boot
Supports Cold boot only

```

This example displays the details when the optical transceiver operates in the muxponder mode.

```

Router#show controllers optics 0/0/0/29 information all
Fri Feb 16 11:06:31.415 UTC

Module State : Ready

Datapath State [Client-0]: Activated

Datapath State [Client-1]: Activated

Datapath State [Client-2]: Activated

Datapath State [Client-3]: Activated

Acquisition Counter:    INVALID

HOST SIDE ALARM COUNTERS
=====
Host-Intf-0-FDD-Alarm-Counter[0]          Host-Intf-0-FED-Alarm-Counter[0]
Host-Intf-1-FDD-Alarm-Counter[0]          Host-Intf-0-FED-Alarm-Counter[0]
Host-Intf-2-FDD-Alarm-Counter[0]          Host-Intf-0-FED-Alarm-Counter[0]
Host-Intf-3-FDD-Alarm-Counter[0]          Host-Intf-0-FED-Alarm-Counter[0]

HOST SIDE FEC-BER FEC-FERC CURRENT VALUES
=====
Host-Intf-0-FEC-BER[0.00E+00]             Host-Intf-0-FEC-FERC [0.00E+00]
Host-Intf-1-FEC-BER[0.00E+00]             Host-Intf-1-FEC-FERC [0.00E+00]
Host-Intf-2-FEC-BER[0.00E+00]             Host-Intf-2-FEC-FERC [0.00E+00]
Host-Intf-3-FEC-BER[0.00E+00]             Host-Intf-3-FEC-FERC [0.00E+00]

Supported Loopback Types :
=====
[1.] Media Internal
[2.] Media Line
[3.] Host Line
[4.] Host Internal

```

[5.] Host Per Lane
 [6.] Media Per Lane
 [7.] Simultaneous Media Host

Display of Alarms for Coherent Optical Transceivers

Table 65: Feature History Table

Feature Name	Release Information	Description
View Additional Alarms for Coherent Optical Transceiver	Release 24.1.1	You can now monitor the optical transceivers for proper functioning and identify the cause of any malfunction. This is made possible because the output of the show alarms command is now enhanced to display the additional media lane alarms for coherent optical transceivers.

Alarms are raised when the optical transceiver malfunctions.

The following media lane alarms are now reported in the **show alarms** command:

- TX loss of alignment - indicates loss of signal alignment on the transmitter.
- TX out of alignment - indicates that the signal on the transmitter (TX) is out of alignment.
- TX CMU loss of lock - indicates that the transmitter (TX) has lost connection (locked) with the external clock (clock monitor unit).
- TX reference clock loss of lock - indicates that the transmitter (TX) has lost connection (locked) with the reference clock.
- TX deskew loss of lock - for traffic flow in the TX direction, this alarm indicates that the end receiver cannot align the physical lanes using alignment marker.
- TX FIFO error - indicates signal FIFO error on the transmitter.
- RX demodulator loss of lock - indicates that the media demodulator cannot achieve lock.
- RX CDC loss of lock - indicates that the receiver (RX) has lost connection (locked) with the external clock.
- RX loss of alignment - indicates that the signal alignment on the receiver (RX) is lost.
- RX out of alignment - indicates that the signal on the receiver (RX) is out of alignment.
- RX deskew loss of lock - for traffic flow in the RX direction, this alarm indicates that the end receiver cannot align the physical lanes using alignment marker.
- RX FIFO error - indicates signal FIFO error on the receiver.
- RX FEC excessive degrade - indicates that the signal has reached or exceeded FED threshold.
- RX FEC detected degrade - indicates that the signal has reached or exceeded FDD threshold.

- Remote degrade - indicates remote signal degradation.
- Local degrade - indicates local signal degradation.
- Remote Phy fault - indicates remote signal fault (RPF).

Restrictions

The optical transceivers may not display these alarms if:

- Optical transceiver is disabled and the optical transceiver is in the **secondary admin-state maintenance** mode.
- Higher priority alarms such as improper removal and loss of signal (LOS) are reported. In such instances, these alarms can be viewed using the **show alarms brief suppressed** command.

Verification

The following **show alarms** command displays the alarms for the coherent optical transceivers.

```
Router#show alarms brief system active
```

```
-----  
Active Alarms  
-----
```

Location	Severity	Group	Set Time	Description
0/RP0/CPU0 Failure With Cisco Licensing Cloud	Major	Software	03/30/2023 12:30:39 UTC	Communications
0/RP0/CPU0 hw_optics: RX POWER LANE-0	Minor	Software	04/02/2023 13:55:38 UTC	Optics0/0/0/31 - HIGH WARNING
0/RP0/CPU0 hw_optics: Optics media rx signal power high warning	Major	Software	04/02/2023 13:55:38 UTC	Optics0/0/0/31 -
0/RP0/CPU0 hw_optics: Optics media tx fifo error	Major	Software	04/02/2023 14:00:01 UTC	Optics0/0/0/29 -
0/RP0/CPU0 hw_optics: Optics media tx loss of alignment	Major	Software	04/02/2023 14:00:01 UTC	Optics0/0/0/29 -
0/RP0/CPU0 hw_optics: Optics media tx out of alignment	Major	Software	04/02/2023 14:00:01 UTC	Optics0/0/0/29 -
0/RP0/CPU0 hw_optics: Optics media tx CMU loss of lock	Major	Software	04/02/2023 14:00:01 UTC	Optics0/0/0/29 -

```
0/RP0/CPU0      Major      Software      04/02/2023 14:00:01 UTC  Optics0/0/0/29 -
hw_optics: Optics media tx reference clock loss of lock

0/RP0/CPU0      Major      Software      04/02/2023 14:00:01 UTC  Optics0/0/0/29 -
hw_optics: Optics media tx deskew loss of lock

0/RP0/CPU0      Major      Software      04/02/2023 14:00:01 UTC  Optics0/0/0/29 -
hw_optics: Optics media rx loss of alignment

0/RP0/CPU0      Major      Software      04/02/2023 14:00:01 UTC  Optics0/0/0/29 -
hw_optics: Optics media rx out of alignment

0/RP0/CPU0      Major      Software      04/02/2023 14:00:01 UTC  Optics0/0/0/29 -
hw_optics: Optics media rx fifo error

0/RP0/CPU0      Major      Software      04/02/2023 14:00:01 UTC  Optics0/0/0/29 -
hw_optics: Optics media rx demodulation loss of lock
```

Loopback on Optical Transceivers

Table 66: Feature History Table

Feature Name	Release Information	Description
Loopback on Optical Transceivers	Release 7.11.1	

Feature Name	Release Information	Description
		<p>You can now easily detect link failures between the optical transceiver and an external device such as a router by creating a loopback within the transceiver itself. Enabling loopback detects the fault in the physical or network connections, such as, traffic loss or a faulty optical transceiver.</p> <p>The loopback configuration allows incoming traffic within the transceiver to be redirected back to its source. By analyzing the loopback signals received at the source, it becomes possible to detect physical connectivity failures or network issues, such as packet loss or a malfunctioning transceiver.</p> <p>The feature introduces these changes:</p> <p>CLI:</p> <p>Modified the controller optics command by adding the following keywords:</p> <ul style="list-style-type: none"> • host loopback internal • host loopback line • loopback internal • loopback line <p>The information loopback keyword is added to the show controller optics command.</p> <p>YANG Data Model:</p> <ul style="list-style-type: none"> • New XPath for <code>Cisco-IOS-XR-controller-optics-cfg.yang</code> <p>(see GitHub, YANG Data Models Navigator)</p>

You can now enable loopback functionality on the optical transceivers. Loopback is the process of redirecting inbound traffic or data signals from an optical transceiver back to the module itself. Re-routing traffic to its source enables utilization of the received data for diagnostic purposes, particularly in the identification and

resolution of physical connectivity issues or network-related problems, such as traffic loss or a faulty optical transceiver.

The optical transceiver is divided into two sides, the host side, which is positioned towards the router, and the media side, which is positioned towards the wire or cable media. It is possible to enable loopback on both the host side and media side of the optical transceiver.



Note Loopback can be performed only when the controller state is active (UP) and in the maintenance mode.

There are four types of loopback:

- Loopback Internal or Media Side Output Loopback
- Loopback Line or Media Side Input Loopback
- Host Loopback Internal or Host Side Input Loopback
- Host Loopback Line or Host Side Output Loopback




Note Configuring the internal loopback brings up the host interface and configuring the line loopback brings up the remote interface.

Media Side Input Loopback Configuration

In loopback line or media side input loopback, the signals received at the media side are looped back to the media side, indicating that the received data on the media is transmitted back to the media, that is, towards the remote interface. This is indicated by the arrow labeled as 1 in the illustration.

Figure 17: Media Side Input Loopback on the Optical Transceiver



Configuration Example

This example shows how to enable media side input loopback on the optical transceiver:

```
Router#config
Router(config)#controller optics 0/0/0/9
Router(config-Optics)#secondary-admin-state maintenance
Router(config-Optics)#loopback line
```



```

Loopback is a traffic-affecting operation
Router(config-Optics)#commit
Router(config-Optics)#end

```

Running Configuration

This example shows the running configuration of the media side input loopback on the optical transceiver:

```

Router#show run controller optics 0/0/0/9
controller Optics0/0/0/9
  secondary-admin-state maintenance
  loopback line
!
```

Verification

This example shows how to verify the media side input loopback configuration on the optical transceiver:

```

Router#show controller optics 0/0/0/9
Controller State: Up
Transport Admin State: Maintenance
Laser State: On
LED State: Green
FEC State: FEC ENABLED
Optics Status:

Optics Type: QSPD 400G FR4
Wavelength: 1301.00 nm
Loopback Host: None
Loopback Media: Line


Alarm Status:
-----
Detected Alarms: None
LOS/LOL/Fault Status:
Performance Monitoring: Disable

```

Media Side Output Loopback

In loopback internal or media side output loopback, the loopback signal originating from the NPU is looped back to the same NPU on the media or line side, towards the remote interface. This is indicated by the arrow labeled as 1 in the illustration.

Figure 18: Media Side Output Loopback on the Optical Transceiver



Configuration Example

This example shows how to enable media side output loopback on the optical transceiver:

```
Router#config
Router(config)#controller optics 0/0/0/9
Router(config-Optics)#secondary-admin-state maintenance
Router(config-Optics)#loopback internal
Loopback is a traffic-affecting operation
Router(config-Optics)#commit
Router(config-Optics)#end
```

Running Configuration

This example shows the running configuration of the media side output loopback on the optical transceiver:

```
Router#show run controller optics 0/0/0/9
controller Optics0/0/0/9
  secondary-admin-state maintenance
  loopback internal
!
```

Verification

This example shows how to verify the media side output loopback configuration on the optical transceiver:

```
Router#show controller optics 0/0/0/9
Controller State: Up
Transport Admin State: Maintenance
Laser State: On
LED State: Green
FEC State: FEC ENABLED
Optics Status:

Optics Type: QSFPDD 400G FR4
Wavelength: 1301.00 nm
Loopback Host: None
Loopback Media: Internal

Alarm Status:
-----
Detected Alarms: None
LOS/LOL/Fault Status:
Performance Monitoring: Disable
```


Host Side Input Loopback Configuration

In host loopback internal or host side input loopback, the loopback signal coming from the NPU is looped back to the NPU on the host, that is, towards the local interface. This is indicated by the arrows labeled as 1 and 2 in the illustration.

Host Side Output Loopback Configuration

In host loopback line or host side output loopback, the signals received at the host side are looped back to the host side, indicating that the received data on the host is transmitted back to the host, that is, towards the local interface. This is indicated by the arrows labeled as 1 and 2 in the illustration.

Figure 20: Host Side Output Loopback on the Optical Transceiver



Configuration Example

This example shows how to enable host side output loopback on the optical transceiver:

```
Router#config
Router(config)#controller optics 0/0/0/9
Router(config-Optics)#secondary-admin-state maintenance
Router(config-Optics)#host loopback internal
Loopback host is a traffic-affecting operation
Router(config-Optics)#commit
Router(config-Optics)#end
```

Running Configuration

This example shows the running configuration on the optical transceiver:

```
Router#show run controller optics 0/0/0/9
controller Optics0/0/0/9
  secondary-admin-state maintenance
  host loopback internal
!
```

Verification

This example shows how to verify the host side output loopback configuration on the optical transceiver:

```
Router#show controller optics 0/0/0/9
Controller State: Up
Transport Admin State: Maintenance
Laser State: On
LED State: Green
FEC State: FEC ENABLED
Optics Status:

Optics Type: QSFPDD 400G FR4
Wavelength: 1301.00 nm
Loopback Host: Internal
Loopback Media: None

Alarm Status:
```

```
-----  
Detected Alarms: None  
LOS/LOL/Fault Status:  
Performance Monitoring: Disable
```




CHAPTER 17

Configuring QDD Optical Line System

This chapter describes the QDD Optical Line System (OLS) and its supported configurations.

- [Overview of QDD Optical Line System, on page 342](#)
- [Supported Wavelength or Frequency Configuration, on page 344](#)
- [Functional Description of QDD OLS, on page 345](#)
- [QDD OLS Configurations, on page 346](#)
- [Use Case for QDD OLS pluggable, on page 356](#)
- [OLS Alarms Troubleshooting, on page 357](#)

Overview of QDD Optical Line System

Table 67: Feature History Table

Feature Name	Release Information	Description
QDD Optical Line System	Release 24.1.1	


Feature Name	Release Information	Description
		<p>The QDD Optical Line System (OLS) is a new pluggable optical amplifier that interconnects two routers or switches for transmitting traffic on a limited number of coherent optical channels over a single span point-to-point link. With the QDD OLS pluggable, it's now possible to obtain the functionality of amplification into a QSFP-DD module that can be plugged into a port of the router or switch. The QDD OLS is supported on Cisco 8201, 8201-32FH, and 8201-24H8FH routers.</p> <p>This feature introduces the following:</p> <ul style="list-style-type: none"> • CLI: <ul style="list-style-type: none"> • controller ots (QDD OLS) • rx-low-threshold • tx-low-threshold • ampli-control-mode • egress-ampli-gain • egress-ampli-power • egress-ampli-safety-control-mode • egress-ampli-osri • show controllers ots (QDD OLS) • YANG Data Model: <ul style="list-style-type: none"> • <code>Cisco-IOS-XR-controller-ots-oper.yang</code> • <code>Cisco-IOS-XR-controller-ots-cfg.yang</code> • <code>Cisco-IOS-XR-pengine-oper.yang</code> • <code>Cisco-IOS-XR-pengine-cfg.yang</code> • <code>Cisco-IOS-XR-pengine-clear-act.yang</code>

The QDD OLS is a pluggable optical amplifier that interconnects two routers or switches for transporting a limited number of coherent optical channels over a single span point-to-point link.

Currently, an optical line system (OLS) requires a separate Cisco Network Convergence System 1000 Series or Cisco Network Convergence System 2000 Series optical system with dedicated 48 or 64 channels MUX/DMX units and amplifiers. The new QDD OLS module provides amplification functionality in a QSFP module, while the passive cables provide MUX/DMX functionality. With this solution, you can compress a point-to-point DWDM system directly into the routing or switching platform.

The new optical line system can now transport 8 or 16 optical channels without any additional optical hardware unit. With the use of external passive MUX/DMX units (examples for these) along with this setup, it's also possible to obtain up to 32 optical channels.

Figure 21: QDD Optical Line System



Supported Wavelength or Frequency Configuration

For each channel supported through ONS-BRK-CS-8LC or ONS-BRK-CS-16LC passive/mux cable, the wavelength or the frequency must be configured according to the table below.


Table 68: QDD OLS Operating Signal Wavelength Range

Channel Spacing	Total Bandwidth	Wavelength		Frequency	
		Start	End	Start	End
8 channels - 200 GHz spaced	19.2 nm 2.4 THz	1539.1 nm	1558.4 nm	192.375 THz	194.775 THz
16 channels - 100 GHz spaced					

Functional Description of QDD OLS

The QDD OLS pluggable contains the COM side and the Line side as shown in the figure below:

Figure 22: Functional Description of QDD OLS



Each physical port of the QDD OLS pluggable is represented as two ots controllers (subport 0 and subport 1). COM port is subport 0 and Line port is subport 1.

The Gain of the Booster is associated to subport 1 while the gain of the Preamplifier is associated to subport 0.

Controller	Optical Ports
ots R/S/I/P/0	COM-RX (booster input)
	COM-TX (preamplifier output)
ots R/S/I/P/1	LINE-RX (preamplifier input)
	LINE-TX (booster output)

QDD OLS Configurations

The following section contains the QDD OLS configuration details.

Configuring the Operational Mode, Amplifier Gain, and Amplifier Output Power

You can configure the mode of operation of the OLS pluggable to either gain control or power control mode.

In the gain control mode, you can configure the desired gain value of the OLS pluggable.

In the power control mode, you can configure the desired output power (TX) of the OLS pluggable.

Gain Control Operational Mode and Amplifier Gain Configuration Example

The following example shows how to configure the gain control operational mode and the amplifier gain of the OLS pluggable:

```
Router#config
Router(config)#controller ots 0/0/2/1/0
Router(config-Ots)#ampli-control-mode manual
Router(config-Ots)#egress-ampli-gain 150
Router(config-Ots)#commit
Router(config-Ots)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration for the OLS pluggable:

```
Router#show run controller ots 0/0/2/1/0
controller Ots0/0/2/1/0
  ampli-control-mode manual
  egress-ampli-gain 150
!
```

Verification

This example shows how to verify the configured gain control operational mode and amplifier gain of the OLS pluggable:

```
Router#show controllers ots 0/0/2/1/0
Thu Mar 23 21:33:49.862 UTC

Controller State: Up

Transport Admin State: In Service

LED State: Green

Alarm Status:
-----
Detected Alarms: None

Alarm Statistics:
-----
RX-LOS-P = 4
RX-LOC = 0
TX-POWER-FAIL-LOW = 1
INGRESS-AUTO-LASER-SHUT = 0
INGRESS-AUTO-POW-RED = 0
```

```

INGRESS-AMPLI-GAIN-LOW = 0
INGRESS-AMPLI-GAIN-HIGH = 0
EGRESS-AUTO-LASER-SHUT = 0
EGRESS-AUTO-POW-RED = 0
EGRESS-AMPLI-GAIN-LOW = 4
EGRESS-AMPLI-GAIN-HIGH = 1
HIGH-TX-BR-PWR = 0
HIGH-RX-BR-PWR = 0
SPAN-TOO-SHORT-TX = 0
SPAN-TOO-SHORT-RX = 0

```

Parameter Statistics:

```

-----
Total Tx Power = 16.72 dBm
Rx Signal Power = -22.29 dBm
Tx Signal Power = 16.53 dBm
Egress Ampli Gain = 14.7 dB
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF

```

Configured Parameters:

```

-----
Egress Ampli Gain = 15.0 dB
Egress Ampli Power = 4.0 dBm
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
Ampli Control mode = Manual
Rx Low Threshold = -30.0 dBm
Tx Low Threshold = -5.0 dBm

```

```

Temperature = 27.92 Celsius
Voltage = 3.33 V

```

Optical Module Details

```

Optics type           : QDD DUAL EDFA
Name                  : CISCO-ACCELINK
OUI Number            : 00.00.00
Part Number           : EDFA-211917-QDD
Rev Number            : 26
Serial Number         : ACW2650Z004
PID                   : ONS-QDD-OLS
VID                   : ves1
Firmware Version      : 2.06
Date Code (yy/mm/dd) : 22/12/31
Fiber Connector Type  : CS

```

Power Control Operational Mode and Amplifier Output Power Configuration Example

The following example shows how to configure the power control operational mode and the amplifier output power of the OLS pluggable :

```

Router#config
Router(config)#controller ots 0/0/2/1/0
Router(config-Ots)#ampli-control-mode powermode
Router(config-Ots)#egress-ampli-power 50
Router(config-Ots)#commit
Router(config-Ots)#exit
Router(config)#exit

```

Running Configuration

This example shows the running configuration for the OLS pluggable:

```
Router(controller-Ots)#show run controller ots 0/0/2/1/0
controller Ots0/0/2/1/0
  ampli-control-mode powermode
  egress-ampli-power 50
!
```

Verification

This example shows how to verify the configured power control operational mode and amplifier output power of the OLS pluggable:

```
Router#show controllers ots 0/0/2/1/0
Thu Jun 1 08:56:37.236 UTC
```

```
Controller State: Up
```

```
Transport Admin State: In Service
```

```
LED State: Green
```

```
Alarm Status:
-----
Detected Alarms: None

Alarm Statistics:
-----
RX-LOS-P = 4
RX-LOC = 0
TX-POWER-FAIL-LOW = 1
INGRESS-AUTO-LASER-SHUT = 0
INGRESS-AUTO-POW-RED = 0
INGRESS-AMPLI-GAIN-LOW = 0
INGRESS-AMPLI-GAIN-HIGH = 0
EGRESS-AUTO-LASER-SHUT = 0
EGRESS-AUTO-POW-RED = 0
EGRESS-AMPLI-GAIN-LOW = 4
EGRESS-AMPLI-GAIN-HIGH = 1
HIGH-TX-BR-PWR = 0
HIGH-RX-BR-PWR = 0
SPAN-TOO-SHORT-TX = 0
SPAN-TOO-SHORT-RX = 0

Parameter Statistics:
-----
Total Tx Power = 5.00 dBm
Rx Signal Power = -22.29 dBm
Tx Signal Power = 4.99 dBm
Egress Ampli Gain = 3.2 dB
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF

Configured Parameters:
-----
Egress Ampli Gain = 15.0 dB
Egress Ampli Power = 5.0 dBm
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
Ampli Control mode = Power
Rx Low Threshold = -30.0 dBm
Tx Low Threshold = -5.0 dBm

Temperature = 29.33 Celsius
Voltage = 3.34 V
```

```
Optical Module Details
```

```

Optics type           : QDD DUAL EDFA
Name                  : CISCO-ACCELINK
OUI Number            : 00.00.00
Part Number           : EDFA-211917-QDD
Rev Number            : 26
Serial Number         : ACW2650Z004
PID                   : ONS-QDD-OLS
VID                   : ves1
Firmware Version      : 2.06
Date Code (yy/mm/dd) : 22/12/31
Fiber Connector Type  : CS

```

Configuring the Low-Threshold Power

You can set the low-threshold power of the optical signal (RX and TX) that can be received or transmitted from the OLS pluggable. The TX and RX low-threshold power values are used for raising and clearing the TX and RX power fail alarms.

Low-Threshold Power Configuration Example

The following example shows how to configure the optical receive (RX) low-threshold power on the OLS pluggable:

```

Router#config
Router(config)#controller ots 0/0/2/1/0
Router(config-Ots)#rx-low-threshold -200
Router(config-Ots)#commit
Router(config-Ots)#exit
Router(config)#exit

```



Note To configure the optical transmit (TX) low-threshold power on the OLS pluggable, use the **tx-low-threshold tx-low** command.

Running Configuration

This example shows the running configuration for the OLS pluggable:

```

Router#show run controller ots 0/0/2/1/0
controller Ots0/0/2/1/0
  rx-low-threshold -200
!
!

```

Verification

This example shows how to verify the configured optical receive (RX) low-threshold power value for the OLS pluggable:

```

Router#show controllers ots 0/0/2/1/0

Controller State: Up

Transport Admin State: In Service

LED State: Yellow

Alarm Status:
-----

```

Detected Alarms:
RX-LOS-P

Alarm Statistics:

```
-----
RX-LOS-P = 0
RX-LOC = 0
TX-POWER-FAIL-LOW = 0
INGRESS-AUTO-LASER-SHUT = 0
INGRESS-AUTO-POW-RED = 0
INGRESS-AMPLI-GAIN-LOW = 0
INGRESS-AMPLI-GAIN-HIGH = 0
EGRESS-AUTO-LASER-SHUT = 0
EGRESS-AUTO-POW-RED = 0
EGRESS-AMPLI-GAIN-LOW = 0
EGRESS-AMPLI-GAIN-HIGH = 0
HIGH-TX-BR-PWR = 0
HIGH-RX-BR-PWR = 0
SPAN-TOO-SHORT-TX = 0
SPAN-TOO-SHORT-RX = 0
```

Parameter Statistics:

```
-----
Total Tx Power = -50.00 dBm
Rx Signal Power = -50.00 dBm
Tx Signal Power = -50.00 dBm
Egress Ampli Gain = 0.0 dB
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
```

Configured Parameters:

```
-----
Egress Ampli Gain = 20.0 dB
Egress Ampli Power = 8.0 dBm
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
Ampli Control mode = Manual
Rx Low Threshold = -20.0 dBm
Tx Low Threshold = -5.0 dBm
```

Temperature = 14.29 Celsius
Voltage = 3.37 V

Optical Module Details

```
Optics type       : QDD DUAL EDFA
Name              : CISCO-ACCELINK
OUI Number       : 00.00.00
Part Number      : EDFA-211917-QDD
Rev Number       : 26
Serial Number    : ACW2650Z004
PID              : ONS-QDD-OLS
VID              : ves1
Firmware Version : 2.06
Date Code (yy/mm/dd) : 22/12/31
Fiber Connector Type : CS
```


Configuring the Optical Safety Remote Interlock (OSRI)

To shut down the amplifier, use the Optical Safety Remote Interlock (OSRI) configuration. This configuration is used during the maintenance of the pluggable, debugging scenarios, and when the OLS pluggable isn't in use. With this configuration enabled, the output power can still be a maximum of -15dBm based on the input power.

OSRI Configuration Example

The following example shows how to configure the Optical Safety Remote Interlock (OSRI) on the OLS pluggable:

```
Router#config
Router(config)#controller ots 0/0/2/1/0
Router(config-Ots)#egress-ampli-osri on
Router(config-Ots)#commit
Router(config-Ots)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration for the OLS pluggable:

```
Router#show run controller optics 0/0/2/1/0
controller Ots0/0/2/1/0
  egress-ampli-osri on
!
```

Verification

This example shows how to verify the configured OSRI for the OLS pluggable:

```
Router#show controllers ots 0/0/2/1/0

Thu Jun  1 09:04:10.335 UTC

Controller State: Up

Transport Admin State: In Service

LED State: Green

Alarm Status:
-----
Detected Alarms: None

Alarm Statistics:
-----
RX-LOS-P = 4
RX-LOC = 0
TX-POWER-FAIL-LOW = 1
INGRESS-AUTO-LASER-SHUT = 0
INGRESS-AUTO-POW-RED = 0
INGRESS-AMPLI-GAIN-LOW = 0
INGRESS-AMPLI-GAIN-HIGH = 0
EGRESS-AUTO-LASER-SHUT = 0
EGRESS-AUTO-POW-RED = 0
EGRESS-AMPLI-GAIN-LOW = 4
EGRESS-AMPLI-GAIN-HIGH = 1
HIGH-TX-BR-PWR = 0
HIGH-RX-BR-PWR = 0
SPAN-TOO-SHORT-TX = 0
SPAN-TOO-SHORT-RX = 0
```

```

Parameter Statistics:
-----
Total Tx Power = -50.00 dBm
Rx Signal Power = -22.36 dBm
Tx Signal Power = -50.00 dBm
Egress Ampli Gain = 0.0 dB
Egress Ampli OSRI = ON
Egress Ampli Force APR = OFF

Configured Parameters:
-----
Egress Ampli Gain = 15.0 dB
Egress Ampli Power = 5.0 dBm
Egress Ampli OSRI = ON
Egress Ampli Force APR = OFF
Ampli Control mode = Power
Rx Low Threshold = -30.0 dBm
Tx Low Threshold = -5.0 dBm

Temperature = 27.90 Celsius
Voltage = 3.34 V

Optical Module Details

Optics type           : QDD DUAL EDFA
Name                  : CISCO-ACCELINK
OUI Number            : 00.00.00
Part Number           : EDFA-211917-QDD
Rev Number            : 26
Serial Number         : ACW2650Z004
PID                   : ONS-QDD-OLS
VID                   : ves1
Firmware Version      : 2.06
Date Code (yy/mm/dd) : 22/12/31
Fiber Connector Type  : CS

```

Configuring Safety Control Mode

You can enable safety control mode only on subport 1.

With safety-control-mode set as **auto** and if LOS is detected on the line RX, the line TX normalizes the signal output power to 8 dBm and the ALS (automatic laser shutdown) and APR (automatic power reduction) alarms are raised.

Safety Control Configuration Example

The following example shows how to enable safety control mode on the OLS pluggable (on subport 1):

```

Router#config
Router (config)#controller ots 0/0/2/1/1
Router (config-Ots)#egress-ampli-safety-control-mode auto
Router (config-Ots)#commit
Router (config-Ots)#exit
Router (config)#exit

```

Running Configuration

This example shows the running configuration for the OLS pluggable:

```

Router#show run controller ots 0/0/2/1/1
controller Ots0/0/2/1/1
  ampli-control-mode manual

```

```
egress-ampli-gain 230
egress-ampli-safety-control-mode auto
!
```

Verification

This example shows how to verify the configured safety control mode:

```
Router#show controllers ots 0/0/2/1/1
```

```
Thu Jun 1 09:04:17.550 UTC
```

```
Controller State: Down
```

```
Transport Admin State: In Service
```

```
LED State: Yellow
```

```
Alarm Status:
```

```
-----
```

```
Detected Alarms:
```

```

RX-LOS-P
EGRESS-AUTO-LASER-SHUT
EGRESS-AUTO-POW-RED
EGRESS-AMPLI-GAIN-HIGH
```

```
Alarm Statistics:
```

```
-----
```

```

RX-LOS-P = 12
RX-LOC = 0
TX-POWER-FAIL-LOW = 1
INGRESS-AUTO-LASER-SHUT = 0
INGRESS-AUTO-POW-RED = 0
INGRESS-AMPLI-GAIN-LOW = 0
INGRESS-AMPLI-GAIN-HIGH = 0
EGRESS-AUTO-LASER-SHUT = 13
EGRESS-AUTO-POW-RED = 13
EGRESS-AMPLI-GAIN-LOW = 2
EGRESS-AMPLI-GAIN-HIGH = 12
HIGH-TX-BR-PWR = 0
HIGH-RX-BR-PWR = 0
SPAN-TOO-SHORT-TX = 0
SPAN-TOO-SHORT-RX = 0
```

```
Parameter Statistics:
```

```
-----
```

```

Total Tx Power = 8.08 dBm
Rx Signal Power = -50.00 dBm
Tx Signal Power = 5.61 dBm
Egress Ampli Gain = 28.9 dB
Egress Ampli Safety Control mode = auto
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
```

```
Configured Parameters:
```

```
-----
```

```

Egress Ampli Gain = 23.0 dB
Egress Ampli Power = 3.0 dBm
Egress Ampli Safety Control mode = auto
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
Ampli Control mode = Manual
Rx Low Threshold = -30.0 dBm
```

```
Tx Low Threshold = -5.0 dBm
```

```
Temperature = 23.00 Celsius
Voltage = 3.36 V
```

Optical Module Details

```
Optics type           : QDD DUAL EDFA
Name                  : CISCO-ACCELINK
OUI Number            : 00.00.00
Part Number           : EDFA-211917-QDD
Rev Number            : 26
Serial Number         : ACW2650Z004
PID                   : ONS-QDD-OLS
VID                   : ves1
Firmware Version     : 2.06
Date Code (yy/mm/dd) : 22/12/31
Fiber Connector Type  : CS
```

Configuring Force APR

Automatic Power Reduction (APR) is an eye-safe output power level of the OLS pluggable when you restart the pluggable (amplifier).

If the OLS pluggable is in the APR state, then the default value for APR power is 8 dBm and APR timer is 9 seconds. When **Force APR** is configured or enabled, the OLS pluggable continues to remain in the APR state. If force APR is not configured or disabled, then the OLS pluggable remains in the Gain mode or Power mode, based on the **ampli-control-mode** user configuration.

Force APR Configuration Example

The following example shows how to configure force APR on the OLS pluggable :

```
Router#config
Router(config)#controller ots 0/0/2/1/1
Router(config-Ots)#egress-ampli-force-apr on
Router(config-Ots)#commit
Router(config-Ots)#exit
Router(config)#exit
```

Running Configuration

This example shows the running configuration for the OLS pluggable:

```
Router#show run controller ots 0/0/2/1/1
controller Ots0/0/2/1/1
  egress-ampli-force-apr on
  ampli-control-mode manual
  egress-ampli-gain 230
  egress-ampli-force-apr on
  egress-ampli-safety-control-mode manual
!
```

Verification

This example shows how to verify the force APR configured:

```
Router#show controllers ots 0/0/2/1/1
```

```
Mon Jan 1 12:37:17.536 UTC
```

```

Controller State: Up

Transport Admin State: In Service

LED State: Green

```

```

Alarm Status:
-----
Detected Alarms: None

Alarm Statistics:
-----
RX-LOS-P = 1
RX-LOC = 0
TX-POWER-FAIL-LOW = 0
INGRESS-AUTO-LASER-SHUT = 0
INGRESS-AUTO-POW-RED = 0
INGRESS-AMPLI-GAIN-LOW = 0
INGRESS-AMPLI-GAIN-HIGH = 0
EGRESS-AUTO-LASER-SHUT = 0
EGRESS-AUTO-POW-RED = 0
EGRESS-AMPLI-GAIN-LOW = 5
EGRESS-AMPLI-GAIN-HIGH = 0
HIGH-TX-BR-PWR = 0
HIGH-RX-BR-PWR = 0
SPAN-TOO-SHORT-TX = 0
SPAN-TOO-SHORT-RX = 0

Parameter Statistics:
-----
Total Tx Power = 17.49 dBm
Rx Signal Power = -19.91 dBm
Tx Signal Power = 17.47 dBm
Egress Ampli Gain = 17.0 dB
Egress Ampli OSRI = OFF
Amplifier Control Mode = Manual
Egress Ampli Force APR = ON

```

```

Configured Parameters:
-----
Egress Ampli Gain = 20.0 dB
Egress Ampli Power = 8.0 dBm
Egress Ampli OSRI = OFF
Egress Ampli Force APR = OFF
Ampli Control mode = Manual
Egress Ampli Force APR = ON
Rx Low Threshold = -30.0 dBm
Tx Low Threshold = -5.0 dBm

```

```

Temperature = 33.40 Celsius
Voltage = 3.25 V

```

Optical Module Details

```

Optics type           : QDD DUAL EDFA
Name                  : CISCO-II-VI
OUI Number            : 00.90.65
Part Number           : 60P310001
Rev Number            : 01
Serial Number         : IIF26410004
PID                   : ONS-QDD-OLS
VID                   : VES1

```

```

Firmware Version      : 1.00
Date Code (yy/mm/dd) : 22/10/10
Fiber Connector Type  : CS

```

Use Case for QDD OLS pluggable

The QDD OLS pluggable can transport 8 or 16 coherent optical channels from the DWDM optical modules that are plugged into the router.

The optical modules are interconnected with the QDD OLS amplifiers using the following cables:

- ONS-BRK-CS-8LC: dual fanout 1x8 cable-assembly with embedded passive splitter and coupler
- ONS-BRK-CS-16LC: dual fanout 1x16 cable-assembly with embedded passive splitter and coupler
- ONS-CAB-CS-LC-5: dual adapter patch-cord CS-connector to LC-connector

The following section explains the 8-channel Optical Line System (OLS) that is achieved by using the QDD OLS pluggable and QDD-400G-ZRP-S modules. With this 8-channel Optical Line System (OLS) set-up it's now possible to obtain 28 dB/112 kilometer span reach. Also, the fiber bandwidth is increased by 8 times.

8-Channel Optical Line System

The following section explains the 8-channel Optical Line System (OLS) that is achieved by using the QDD OLS pluggable and QDD-400G-ZR-S or QDD-400G-ZRP-S module. With this 8-channel Optical Line System (OLS) set up it's now possible to obtain 28 dB/112 kilometer span reach. Also, the fiber bandwidth is increased by 8 times.

This section explains the 8-channel optical line system (OLS) that is achieved by using the following:

- Four Cisco 8000 series router (represented as Node A, Node B, Node C, and Node D)
- Sixteen QDD-400G-ZR-S or QDD-400G-ZRP-S modules
- Two QDD OLS (ONS-QDD-OLS) pluggables
- Two ONS-BRK-CS-8LC breakout cables
- Two ONS-CAB-CS-LC-5 fiber optic cable

Connections on Node A

Four QDD-400G-ZR-S or QDD-400G-ZRP-S modules are inserted into port 0, port 1, port 2, and port 3 of the Cisco 8000 series router. The QDD OLS (ONS-QDD-OLS) pluggable is inserted into port 4 of the Cisco 8000 series router.

Connections on Node B

Four QDD-400G-ZR-S or QDD-400G-ZRP-S modules are inserted into port 0, port 1, port 2, and port 3 of the Cisco 8000 series router.

Connections between Node A and Node B

Using the ONS-BRK-CS-8LC breakout cable connect eight QDD-400G-ZR-S or QDD-400G-ZRP-S modules (four each on Node A and Node B) and the QDD OLS (ONS-QDD-OLS) pluggable (port 3 of Node A).

Connections on Node C

Four QDD-400G-ZR-S or QDD-400G-ZRP-S modules are inserted into port 0, port 1, port 2, and port 3 of the Cisco 8000 series router.

Connections on Node D

Four QDD-400G-ZR-S or QDD-400G-ZRP-S modules are inserted into port 0, port 1, port 2, and port 3 of the Cisco 8000 series router. The QDD OLS (ONS-QDD-OLS) pluggable is inserted into port 4 of the Cisco 8000 series router.

Connections between Node C and Node D

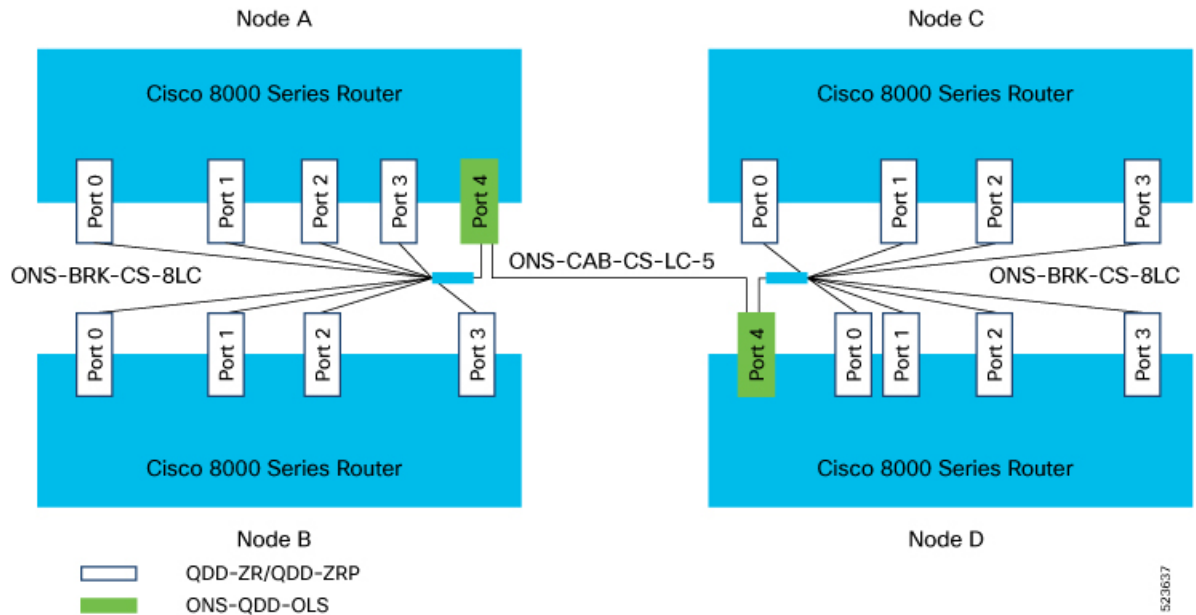
Using the ONS-BRK-CS-8LC breakout cable connect eight QDD-400G-ZR-S or QDD-400G-ZRP-S modules (four each on Node A and Node B) and the QDD OLS (ONS-QDD-OLS) pluggable (port 3 of Node A).

Connections between Node A and Node D

Using the ONS-CAB-CS-LC-5 fiber optic cable connect both the QDD OLS (ONS-QDD-OLS) pluggables that are present in Node A and Node D.

The representation of these interconnections between Node A/Node B and Node C/Node D are depicted in the block diagram below:

Figure 23: 8-Channel Optical Line System



523637

OLS Alarms Troubleshooting

This section contains the procedures for troubleshooting alarms.

RX-LOS-P

Default Severity: Critical

Logical Object: Controller

The RX-LOS-P alarm is raised when there is loss of signal.

Clear the RX-LOS-P Alarm

1. Verify the transmission (TX) at the peer end.
2. Check the fiber connections.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

RX-POWER-FAIL-LOW

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The RX-POWER-FAIL-LOW alarm is raised when the RX power is below the configured low threshold values.

Clear the RX-POWER-FAIL-LOW Alarm

1. Verify the transmission (TX) at the peer end.
2. Check the fiber connections.
3. Increase the peer end gain or transmit-power value to obtain the RX power above the threshold.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

TX-POWER-FAIL-LOW

Default Severity: Critical

Logical Object: Controller

The TX-POWER-FAIL-LOW alarm is raised when the TX power is below the configured low threshold values.

Clear the TX-POWER-FAIL-LOW Alarm

1. Increase the gain or power configuration value to obtain the TX power above the threshold.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

EGRESS-AMPLI-GAIN-LOW

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The EGRESS-AMPLI-GAIN-LOW alarm is raised when the actual gain of the OLS pluggable is lower than the configured gain value.

Clear the EGRESS-AMPLI-GAIN-LOW Alarm

1. Configure the gain value within the optimum range.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

EGRESS-AMPLI-GAIN-HIGH

Default Severity: Minor (MN), Non-Service-Affecting (NSA)

Logical Object: Controller

The EGRESS-AMPLI-GAIN-HIGH alarm is raised when the actual gain of the OLS pluggable is higher than the configured gain value.

Clear the EGRESS-AMPLI-GAIN-HIGH Alarm

1. Verify the RX and TX values and adjust the gain within the optimum working range.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

EGRESS-AUTO-LASER-SHUT

Default Severity: Not-Alarmed

Logical Object: Controller

The EGRESS-AUTO-LASER-SHUT alarm is raised when there is loss of signal (LOS) on the OTS line side (subport 1)

Clear the EGRESS-AUTO-LASER-SHUT Alarm

1. Verify the fiber connections on the line side of the OLS pluggable.
2. Verify the gain or power on the line side of the peer end.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).

EGRESS-AUTO-POW-RED

Default Severity: Not-Alarmed

Logical Object: Controller

The EGRESS-AUTO-POW-RED alarm is raised when there is loss of signal (LOS) on the OTS line side (subport 1)

Clear the EGRESS-AUTO-POW-RED Alarm

1. Verify the fiber connections on the line side of the OLS pluggable.
2. Verify the gain or power on the line side of the peer end.

If the alarm does not clear, log into the Technical Support Website at <http://www.cisco.com/c/en/us/support/index.html> for more information or call Cisco TAC (1 800 553-2447).