



Configuration Guide for Cisco NCS 1004, IOS XR Release 7.2.1

First Published: 2020-07-16

Last Modified: 2026-05-07

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)

© 2020 Cisco Systems, Inc. All rights reserved.



CONTENTS

CHAPTER 1	New and Changed Information	1
	New and Changed Information	1

CHAPTER 2	Cisco NCS 1004 Overview	7
	Cisco NCS 1004 units	7
	Interoperability constraints and timing requirements for Cisco NCS 1001 and Cisco NCS 1004	8
	Supported line cards	8
	1.2T card interoperability with OTN-XP card	10
	Prerequisites for interoperability with OTN-XP card	10
	Scenario on interoperability with OTN-XP card	10
	Configure interoperability between the 1.2T card and the OTN-XP card	11

CHAPTER 3	Configuring the Card Mode	13
	1.2T and 1.2TL Line Cards	13
	Card modes	13
	Sub 50G configuration	14
	Supported Data Rates	15
	Configuring the Card Mode	18
	Regeneration Mode	22
	Configuring the Card in Regen Mode	22
	Verifying the Regen Mode	23
	Configuring the BPS	23
	Configuring the Trunk Rate for BPSK	24
	Viewing the BPSK Trunk Rate Ranges	24
	OTN-XP Card	26
	LC Mode on OTN-XP Card	26

Configuring the LC Mode 27
 Muxponder Configuration on OTN-XP Card 32
 Configuring the Muxponder Mode for 10G Grey Muxponder 33

CHAPTER 4

Configuring Controllers 35

AINS 35
 AINS States 36
 Soak Time Period 36
 Configuring AINS 36
 Disabling AINS 37
 Displaying the AINS Configuration 37
 Configuring AINS on OTN-XP Card 41

FEC 44
 FEC States for Ethernet Controller 44
 Configuring FEC on the Ethernet Controller 46
 FEC States for CoherentDSP Controller 47
 Configuring FEC on CoherentDSP Controllers 47
 Verifying FEC on CoherentDSP Controllers 48

Laser Squelching 49
 Configuring Laser Squelching on OTN-XP Card 50

Idle Insertion 51
 Enabling Idle Insertion on OTN-XP Card 52

LLDP Drop 53
 Configuring LLDP Drop 54
 Verifying the Status of LLDP Drop 54

Link Layer Discovery Protocol (LLDP) Support on Management Interface 56

MAC Address Snooping on Client Ports 60
 Configuring MAC Address Snooping on Client Ports 60
 Viewing Neighbor MAC Address 61

Loopback 62

Restore Factory Settings 68

Headless Mode 69

Trail Trace Identifier 69

Chromatic dispersion 71

Frequency	73
Pseudo Random Binary Sequence	73
Configuring Pseudo Random Binary Sequence	74
Verifying PRBS	75
Viewing PRBS Performance Monitoring Parameters	75
Configuring PRBS on OTN-XP Card	76
Verifying PRBS on OTN-XP Card	77

CHAPTER 5
Performance Monitoring 79

Performance monitoring parameters	79
Instantaneous Q-margins	79
Configure PM parameters	80
View PM parameters	81
Performance monitoring parameter outputs for Ethernet, FEC, optics, and Fibre Channel controllers	83
Clear PM parameters	89

CHAPTER 6
IP Access Lists 91

IP access control lists	91
Best practices and requirements for configuring IP access control lists	92
Configure an IP ACL	92
Verify access control lists	93

CHAPTER 7
Layer 1 Encryption 95

IKEv2 Overview	96
OTNSec Encryption Overview	98
Prerequisites	99
Limitations	100
Configuration Workflow	100
Configuring an IKEv2 Proposal	101
Configuring an IKEv2 Policy	102
Configuring a Keyring	103
Configuring a IKEv2 Profile	104
Configuring an OTNSec Policy	104

Configuring the GCC Interface	105
Configuring OTNSec on ODU4 Controllers	106
Configuration Example	107
Verification	110
Troubleshooting	110
IKEv2 Certificate-Based Authentication	111
Configuring IKEv2 Certificate-Based Authentication	111
You May Be Interested In	115

CHAPTER 8**GMPLS UNI for Packet and Optical Integration 117**

Understanding GMPLS UNI	117
Use Case Overview	118
Prerequisites	118
Limitations	119
Configuration Workflow	119
Configure an LMP link and Alien Wavelength on an NCS 2000 node using CTC	119
Configure an LMP link with Alien Wavelength on NCS 2000 using CTC for signaled unnumbered circuits	122
Retrieve Ifindex from NCS 2000 Node	124
Configure LMP on Cisco NCS 1004 Node	125
Configure RSVP on NCS 1004 Node	126
Configure MPLS Tunnel on a NCS 1004 Node for Numbered Circuit	127
Verification	128
General Troubleshooting	135
You May Be Also Interested In	135

CHAPTER 9**Remote Node Management 137**

Remote node management using GCC	137
Supported and unsupported features of the GCC interface	138
Supported protocols	139
Enable the GCC interface	139
Configure the GCC interface	140
Configure static routes over the GCC interface	141
Configure OSPF routes over the GCC interface	142

iBGP supports over GCC interfaces	143
Caution: Follow restrictions for iBGP support using GCC	144
Enabling the GCC Interface	144
Configuring the Management Interface	144
Configuring the Loopback Interface	145
Configuring the GCC interface	145
Verifying iBGP Support Using GCC	146
iBGP configuration parameters for GCC interfaces	147

CHAPTER 10**Smart Licensing 151**

Understanding Smart Licensing	151
Benefits of Smart Licensing	154
Licensing in NCS 1004	154
Software Entitlements of Cisco NCS 1004	154
Configure Smart Licensing	155
Creating a Token	156
Verifying Smart Licensing Configuration	157
License Registration	160
Smart Licensing for OTN-XP Line Card	161
Checking the License Usage Count	161

CHAPTER 11**USB Device Automount 165**

USB automounts	165
Mount a USB device in virtual machine environments	165
Unmount a USB device	166

CHAPTER 12**Fault Profiles 169**

Fault profiles	169
Best practice for configuring fault profiles	170
Configure a fault profile	170

APPENDIX A**Supported SNMP MIBs on NCS 1004 171**



CHAPTER 1

New and Changed Information

This chapter lists the new and changed information for each release.

- [New and Changed Information, on page 1](#)

New and Changed Information

See [Data Models Configuration Guide for Cisco NCS 1004](#) and [Telemetry Configuration Guide for Cisco NCS 1000 Series](#) to refer the other configuration guides of NCS 1004.

This table summarizes new and changed information for Release 7.2.1, and lists where the features are documented.

Table 1: New and Changed Features - R7.2.1

Feature	Description	Where Documented
Mixed client rate support	In muxponder slice mode, both the slices can be configured with different client rates. For example, slice 0 can be configured with 100GE client rate and slice 1 can be configured with OTU4 client rate and the other way round. This provides flexibility for the card to simultaneously carry both the OTN and Ethernet client traffic across the two slices.	Supported Data Rates and Configuring the Card Mode

Feature	Description	Where Documented
IKEv2 Certificate-based authentication	<p>IKEv2 uses RSA digital signatures to authenticate peer devices before setting up SAs. RSA signatures employ a PKI-based method of authentication.</p> <p>In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other.</p>	Layer 1 Encryption
MAC Address Snooping on Client Ports	MAC address snooping allows you to learn the MAC address of the neighbor, that is connected to the client ports. You can enable ARP snooping on all client ports and learn the MAC address of neighbors through CLI.	MAC Address Snooping on Client Ports
QSFP-28 FR 100G and DAC pluggable support	<p>The following pluggables are supported from Release 7.2.1</p> <ul style="list-style-type: none"> • QSFP-100G-CU1M • QSFP-100G-CU2M • QSFP-100G-CU3M • QSFP-100G-CU5M • QSFP28-100G-FR 	QSFP-28 FR 100G and DAC pluggable support
FEC Mode Support for CoherentDSP Controller	The FEC states can be configured for the CoherentDSP Controller. The supported FEC states are EnhancedSD15 and EnhancedSD27 (default).	Configuring FEC on CoherentDSP Controllers
BPSK Modulation Support	Binary Phase Shift Keying (BPSK) feature enables you to configure the trunk rates using CLI, NetConf YANG, and Open Config (OC) models. The supported trunk rates for the BPSK modulation are 50G, 100G, and 150G.	Configuring the Trunk Rate for BPSK

Feature	Description	Where Documented
iBGP support over GCC	<p>The Internal BGP (iBGP) support over GCC allows external devices to exchange BGP routes through management interfaces of NCS1004 system.</p> <p>The iBGP over GCC feature enables you to configure VPN routing and forwarding (VRF) on the GCC management interfaces (port 0 and port1) of the NCS 1004 device. The VRF enables traffic isolation between the management ports (port 0 and port1).</p>	iBGP Support Using GCC
1.2T Card Interoperability with the NCS1K4-OTN-XP Card	<p>NCS 1004 supports the NCS1K4-OTN-XP card with 100G grey-optics support. The OTN-XP card can be interoperable with the 1.2 Tbps card. In an interoperability scenario, the 1.2 T card can serve as a client port and the OTN-XP card can serve as a trunk port. The trunk port can converge 10 x 10 G traffic and transmit as 100G traffic in the OTU4 mode. This OTU4 traffic can further be multiplexed to a higher bandwidth Dense Wavelength-Division Multiplexing (DWDM) signal by connecting to the 1.2 T OTU4 client interface.</p>	1.2T Card Interoperability with OTN-XP Card
OTN-XP card	<p>The OTN-XP line card supports up to 1.6Tbps of OTN aggregation switching functionality to optimize the available bandwidth. A single line card supports 8x100GE muxponder or 2x400 GE transponder applications.</p> <p>The OTN-XP line card contains:</p> <ul style="list-style-type: none"> • Eight QSFP 28 ports • Four QSFP-DD ports • Two CFP2 ports 	Hardware Installation Guide for Cisco NCS 1004

Feature	Description	Where Documented
LC Mode Configuration on OTN-XP Card	When an OTN-XP card is installed in the NCS 1004 chassis, the card is in POWERED_ON state. A datapath mode must be configured using the LC mode CLI after which the card transitions to the OPERATIONAL state.	LC Mode on OTN-XP Card
Muxponder Configuration on OTN-XP Card	The muxponder configuration on the OTN-XP card supports two slices, 0 and 1. You can configure mxponder-slice 0, mxponder-slice 1, or both. Each mxponder-slice supports 10 client interfaces.	Muxponder Configuration on OTN-XP Card
AINS Configuration on OTN-XP Card	The default AINS settings for all controllers on the OTN-XP card can be configured using the shared plane configuration CLI. However, it is possible to override the default AINS settings on a specific controller using the CLI.	Configuring AINS on OTN-XP Card
Smart Licensing for OTN-XP Line Card	<p>Smart Licensing is a cloud-based, software license management solution that enables you to automate time-consuming, manual licensing tasks.</p> <ul style="list-style-type: none"> • The license calculation is based on 100G client bandwidth and is independent of the client type. • The licensed OTN-XP Line Card PID is NCS1K4-OTN-XPL. • The license is charged based on the usage of 100G client bandwidth. 	Smart Licensing for OTN-XP Line Card

Feature	Description	Where Documented
GCC0 support on OTU Interfaces for the OTN-XP card	The OTN-XP line card provides OTU interface that supports communication channels between adjacent network elements or nodes using GCC bytes in the OTN header. Remote node management is supported over the GCC interface. The node supports GCC0 on corresponding OTU2, OTU2e, and OTU4 interfaces. The node (Cisco FPGA) supports a maximum of 22 GCC channels for each card.	Remote Node Management on OTN-XP Card
Laser Squelching Support on OTN-XP Card	The laser squelching feature when enabled on the 10GE controllers, allows the laser to shut down in the event of trunk faults. The SQUELCHED alarm is raised.	Configuring Laser Squelching on OTN-XP Card
Idle insertion Support on OTN-XP Card	The idle insertion feature enables you to hold the propagation of local faults on the trunk port of the ethernet controller. You can enable the feature by configuring the hold-off timer on the 10GE controllers.	Enabling Idle Insertion on OTN-XP Card
Loopback feature support on Ethernet and OTU Controller	The Loopback feature enables you to configure internal and line loopbacks on the OTU2, OTU2e, OTU4, and 10GE controllers.	Configuring Loopback on OTN-XP Card
PRBS Support on ODU2e Controller	<p>Pseudo Random Binary Sequence (PRBS) feature enables you to perform data integrity checks between the NCS1004 trunk links without enabling the actual client traffic.</p> <p>This feature enables you to configure Optical Channel Payload Unit (OPU) on the ODU2e controller followed by the PRBS mode and the pattern. The PRBS supported pattern on the OTN XP line card is invertedPN31.</p>	Configuring PRBS on OTN-XP Card



CHAPTER 2

Cisco NCS 1004 Overview

This chapter provides an overview of the Cisco Network Convergence Series (NCS) 1004.

- [Cisco NCS 1004 units, on page 7](#)
- [Interoperability constraints and timing requirements for Cisco NCS 1001 and Cisco NCS 1004, on page 8](#)
- [Supported line cards, on page 8](#)
- [1.2T card interoperability with OTN-XP card, on page 10](#)

Cisco NCS 1004 units

A Cisco NCS 1004 unit is an optical transport platform that

- supports up to 4.8 Tbps traffic in a compact two-RU form factor,
- includes two redundant field-replaceable AC and DC power supply units; three redundant field-replaceable fans; a field-replaceable controller card; and SSD disks both on board the chassis and on the controller card. This configuration provides enhanced resiliency, and
- features four line card slots per chassis, each capable of hosting a line card for scalable network expansion.

See [Supported Line Cards](#), for more information.

Key features and benefits of Cisco NCS 1004

The NCS 1004 delivers these benefits:

- Transport of any trunk rate from 150 200 to 600 Gbps wavelengths in 50 Gbps increments on the same platform through software provisioning.
- Support of granular control of baud-rate and modulation format to maximize spectral efficiency.
- One universal transponder optimized for metro, long-haul, and submarine applications.
- Support for up to 350,000 ps/nm of residual chromatic dispersion compensation.
- Ability to transport of 100GE and OTU4 client rates on the same platform through software provisioning.
- 600G DWDM, providing unparalleled scale and density—64 channels of 600G at 75 GHz, delivering 38.4 Tbps in 16 RU.
- State-of-the-art AES-256 encryption at scale—4.8 Tbps of encrypted trunk capacity per 2 RU.

Interoperability constraints and timing requirements for Cisco NCS 1001 and Cisco NCS 1004

Understand the interoperability constraints and timing requirements between Cisco NCS 1001 and Cisco NCS 1004.

Interoperability constraints and requirements

- Traffic loss may occur under certain conditions when Cisco NCS 1001 with Protection Switching Module (PSM) is configured as non-revertive and operates with Cisco NCS 1004.
- If traffic switches from the working path to the protect path, users must wait 120 seconds before performing a manual switch to prevent traffic loss.
- A manual switch performed before the interval, followed by a protect path failure, can result in up to 13 seconds of traffic loss.

Timing limitations

Supported line cards

A supported line card is a hardware module for the Cisco NCS 1004 platform that

- provides specific types and numbers of high-speed client interfaces,
- supports variable trunk port configurations for flexible bandwidth and transmission rates, and
- offers distinctive modulation, frequency, and error-correction capabilities tailored to deployment needs.

These line cards are supported on Cisco NCS 1004.

NCS1K4-1.2T-K9 C-band line card

The NCS1K4-1.2T-K9 (or 1.2 Tbps) C-band line card contains 12 QSFP-28 clients and two DWDM trunk ports. These trunk ports support multiple line rates and offer precise control of the modulation format, baud rate, and forward error correction. The trunk ports can be configured using software. The line card supports both module and slice configurations.



Note "1.2TC" refers to the NCS1K4-1.2T-K9 C-band line card.

The features of the 1.2T line card are:

- The card provides up to 12 100G or OTU4 client ports.
- The baud rate can be controlled between 28 Gbd/s and 72 Gbd/s.
- The frequency range is 191.25 to 196.1 THz with a default value of 193.1 THz.
- The modulation formats can be QPSK, 8 QAM, 16 QAM, 32 QAM, or 64 QAM.

- Configure hybrid modulation formats using 1/128 bits per symbol granularity.
- Forward Error Correction (FEC) of 27% and 15% overhead across line rates (only 15% for 600G).
- In Release 7.1.1, the trunk line rate can be configured from 150G to 600G in 50G increments.
- In Release 7.2.1 and later releases, the trunk line rate can be configured from 50G to 600G in 50G increments.

NCS1K4-1.2TL-K9 L-band line card

The NCS1K4-1.2TL-K9 (or 800 Gbps) L-band line card has 12 QSFP-28 based clients and two DWDM trunk ports. The trunk ports are capable of several line rates with fine control of modulation formats, baud rate, and forward error correction and are software configurable. The line card supports module and slice configurations.



Note "1.2TL" refers to the NCS1K4-1.2TL-K9 L-band line card.



Note There is no support for GMPLS, remote management using GCC, and smart licensing.

The features of the 1.2TL line card are:

- The card provides up to eight 100G or OTU4 client ports.
- The client ports map to two trunk ports that operate on any rate between 200G and 400G with 50G increments.
- The modulation formats can be controlled between QPSK, 8 QAM, and 16 QAM.
- The baud rate can be controlled between 31.5Gbd/s and 72Gbd/s.
- The frequency range is 186.10 to 190.85 THz with a default value of 188.50 THz. Only 100 MHz spacing is supported.
- Hybrid modulations formats can be configured through 1/128 bits/symbol granularity.
- Forward Error Correction (FEC) supports 15% and 27% overhead.

NCS1K4-OTN-XP line card

From R7.2.1 onwards, NCS 1004 supports the NCS1K4-OTN-XP card with 100G grey-optics support.



Note "OTN-XP" refers to the NCS1K4-OTN-XP line card.

The OTN-XP card contains:

- Eight QSFP 28 ports
- Four QSFP-DD ports
- Two CFP2 ports

The OTN-XP card supports up to 1.6Tbps of OTN aggregation switching functionality to optimize the available bandwidth. A single line card supports 8x100GE muxponder or 2x400 GE transponder applications. The OTN-XP card supports 400GE/OTUC4, 100GE/OTU4, 10GE/OTU2/OTU2e, 16G FC, and 32G FC client rates.

For more information on the mode configuration, see [Muxponder Configuration on OTN-XP Card, on page 32](#).

1.2T card interoperability with OTN-XP card

A card interoperability is a network capability that

- enables different types of optical interface cards to exchange and transport data,
- allows client cards to aggregate and convert multiple traffic streams for trunk transmission, and
- requires compatible pluggable optics to ensure reliable communication between platforms.

The OTN-XP card can be interoperable with the 1.2T card. In an interoperability scenario, the 1.2T card can serve as a trunk port and the OTN-XP card can serve as a client port. The trunk port of OTN-XP can converge 10 x 10 G traffic and transmit as 100G traffic in the OTU4 mode. This OTU4 traffic can further be multiplexed to a higher bandwidth DWDM signal by connecting to 1.2T OTU4 client interface.

Additional reference information

For interoperability between the 1.2T and OTN-XP cards, supported pluggable optics are required:

- **On the trunk side (OTN-XP card):** Cisco QSFP-100G-LR4 Pluggable Optics Module (ONS-QSFP28-LR4). The same module should be used on the client side of the 1.2T card to ensure compatibility.
- **On the client side (OTN-XP card):**
 - ONS-QSFP-4x10-MLR
 - QSFP-40G-SR4

Prerequisites for interoperability with OTN-XP card

- Configure the OTN-XP card in the 10x10G traffic mode.
- Configure the 1.2T card in the OTU4 client mode with supported trunk rate.
- Ensure that the software installed on both route processors and cards are stable in the supported traffic modes for the 1.2T and OTN-XP card.


Scenario on interoperability with OTN-XP card

Consider a topology in which the OTN-XP card is configured in 10x10G traffic mode and the 1.2T card is configured in OTU4 client mode with a supported trunk rate.

For the solution to work, the OTN-XP trunk optics must interoperate seamlessly with the 1.2T client optics.

Interoperability scenario

Figure 1: Interoperability Topology



Configure interoperability between the 1.2T card and the OTN-XP card

Establish communication and feature compatibility between the 1.2T card and OTN-XP card.

This procedure is required when you need two different card types—1.2T card and OTN-XP card—to operate together in mixed environments, ensuring optimal performance and interoperability.

Before you begin

- Ensure both cards (1.2T and OTN-XP) are properly installed in your system.
- Confirm that the system software supports the required features and modes.

Follow these steps to configure interoperability between the 1.2T card and the OTN-XP card.

Procedure

-
- Step 1** Configure the muxponder mode on the 1.2T card, see [Configuring the Card Mode](#), on page 18.
- Step 2** Configure the LC mode on the OTN-XP card, see [Configure LC mode](#).
- Step 3** Configure the OTN-XP card in muxponder mode, [Configuring the Muxponder Mode for 10G Grey Muxponder](#), on page 33.
- Step 4** Run the `no shut` command on both trunk ports.

Example:

The is a sample to perform no shut on the trunk port.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller Optics 0/1/0/0
RP/0/RP0/CPU0:ios(config-Optics)#no shut
RP/0/RP0/CPU0:ios(config-Optics)#description trunk port
RP/0/RP0/CPU0:ios(config-Optics)#exit
RP/0/RP0/CPU0:ios(config)#exit
```

Note

Ensure that there are no alarms generated and link recovery after events such as fiber cut, card reload, pluggable Online Insertion and Removal (OIR), and re-provisions.

Step 5 Optionally, perform additional controller configurations on the OTN-XP card's grey optics and 1.2T OTU4 client optics.

- Loopback
- Trail Trace Identifier (TTI)
- Maintenance
- Performance Monitoring (PM) enable
- Threshold

Interoperability between the 1.2T card and OTN-XP card is established with trunk ports enabled and optional controller features configured as needed.



CHAPTER 3

Configuring the Card Mode

This chapter lists the supported configurations and the procedures to configure the card mode on the line cards.



Note Unless otherwise specified, “line cards” refers to 1.2T and 1.2TL line cards.

- [1.2T and 1.2TL Line Cards, on page 13](#)
- [OTN-XP Card, on page 26](#)

1.2T and 1.2TL Line Cards

The following section describes the supported configurations and procedures to configure the card modes on the line cards.

Card modes

A card mode is a configuration option for line cards that

- determines which trunk and client ports operate together,
- controls whether trunk rates are set globally or independently, and
- defines how client ports are mapped to trunk ports.

Each line card includes two trunk ports (0 and 1) and twelve client ports (2 through 13). You can configure each line card in one of two modes.

Muxponder mode

Both trunk ports (0 and 1) are configured with the same trunk rate. Client-to-trunk mapping is sequential and enables all client ports to use the combined bandwidth of the trunks efficiently.

Muxponder slice mode

Each trunk port (0 and 1) is configured independently, allowing different trunk rates for each. The client to trunk mapping is fixed. Client ports 2 through 7 map to trunk 0, and client ports 8 through 13 map to trunk 1.

Sub 50G configuration

A Sub 50G configuration is a line card setup that

- enables data transmission rates below 50 Gbps or coupled mode on muxponder line cards,
- allows flexible port assignments and mapping for various supported data rates, and
- imposes specific requirements and operational restrictions when configured.

Standard port configurations for Sub 50G data rates

You can configure Sub 50G or coupled mode on the line card only when operating in muxponder mode. This table shows the supported port configurations for various data rates.

Table 2: Port configuration for supported data rates

Trunk data rate (per trunk)	Total configured data rate	Card support	Trunk ports	Client ports for trunk 0 (100G)	Shared client port (50G per trunk)	Client ports for trunk 1 (100G)
50G	100G	1.2T, 1.2TL	0, 1	-	2	-
150G	300G	1.2T, 1.2TL	0, 1	2	3	4
250G	500G	1.2T, 1.2TL	0, 1	2, 3	4	5, 6
350G	700G	1.2T, 1.2TL	0, 1	2, 3, 4	5	6, 7, 8
450G	900G	1.2T	0, 1	2, 3, 4, 5	6	7, 8, 9, 10
550G	1.1T	1.2T	0, 1	2, 3, 4, 5, 6	7	8, 9, 10, 11, 12

Alternate port configurations (split client port mapping)

From Release 7.5.2, 1.2T and 1.2TL line cards support an alternate port configuration for Sub 50G (split client port mapping), which you configure using CLI. This table shows the alternate port mapping for the supported data rates:

Operational considerations

In all x50G configurations, client traffic on the middle port is affected by ODUK-BDI and LF alarms after a **power cycle** or **link flap** on the trunk side. This issue occurs when two network lanes operate in coupled mode and move from low to high power. To resolve this, create a new frame at the near-end or far-end by performing a **shut** or **no shut** of the trunk ports.

Restrictions for coupled mode configurations

These restrictions apply to coupled mode configuration:

- Both trunk ports must be configured with the same bits-per-symbol or baud rate and must be sent over the same fiber and direction.
- Chromatic dispersion values must be configured identically for both trunk ports.
- When trunk internal loopback is configured, it must be set for both trunk ports. Configuring internal loopback on only one trunk results in traffic loss.
- A fault on a trunk port of a coupled pair may cause errors on all clients, including those running only on the unaffected trunk port.

Supported Data Rates

The following data rates are supported on the line card.

In R7.0.1, you can configure the client port to OTU4 only in the muxponder mode. In R7.1.1 and later releases, you can configure the client port to OTU4 in both the muxponder and muxponder slice modes. In muxponder slice mode, both the slices must be configured with either OTU4 or 100GE Ethernet client rates in R7.1.1. In R7.2.0, a mixed configuration of OTU4 and 100GE is supported in the muxponder slice mode. LLDP drop, L1 encryption, and AINS are not supported on the OTU4 configuration.

The following table displays the client and trunk ports that are enabled for the muxponder configuration.

Trunk Data Rate	Card Support	Client Data Rate (100GE, OTU4)	Trunk Ports	Client Ports
100	1.2T, 1.2TL	100GE, OTU4	0, 1	2, 3
200	1.2T, 1.2TL	100GE, OTU4	0, 1	2, 3, 4, 5
300	1.2T, 1.2TL	100GE, OTU4	0, 1	2, 3, 4, 5, 6, 7
400	1.2T, 1.2TL	100GE, OTU4	0, 1	2, 3, 4, 5, 6, 7, 8, 9
500	1.2T	100GE, OTU4	0, 1	2, 3, 4, 5, 6, 7, 8, 9, 10, 11
600	1.2T	100GE, OTU4	0, 1	2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13

The following table displays the client and trunk ports that are enabled for the muxponder slice 0 configuration.

Trunk Data Rate	Card Support	Client Data Rate	Trunk Ports	Client Ports
100	1.2T, 1.2TL	100GE, OTU4	0	2
200	1.2T, 1.2TL	100GE, OTU4	0	2, 3
300	1.2T, 1.2TL	100GE, OTU4	0	2, 3, 4

Trunk Data Rate	Card Support	Client Data Rate	Trunk Ports	Client Ports
400	1.2T, 1.2TL	100GE, OTU4	0	2, 3, 4, 5
500	1.2T	100GE, OTU4	0	2, 3, 4, 5, 6
600	1.2T	100GE, OTU4	0	2, 3, 4, 5, 6, 7

The following table displays the client and trunk ports that are enabled for the muxponder slice 1 configuration.

Trunk Data Rate	Card Support	Client Data Rate	Trunk Ports	Client Ports
100	1.2T, 1.2TL	100GE, OTU4	1	8
200	1.2T, 1.2TL	100GE, OTU4	1	8, 9
300	1.2T, 1.2TL	100GE, OTU4	1	8, 9, 10
400	1.2T, 1.2TL	100GE, OTU4	1	8, 9, 10, 11
500	1.2T	100GE, OTU4	1	8, 9, 10, 11, 12
600	1.2T	100GE, OTU4	1	8, 9, 10, 11, 12, 13

All configurations can be accomplished by using appropriate values for client bitrate and trunk bitrate parameters of the **hw-module** command.

The following table displays the trunk parameter ranges for the 1.2T card.

Trunk Payload	FEC	Min BPS	Max BPS	Min GBd	Max GBd
50G	15%	1	1.3125	24.0207911	31.5272884
50G	27%	1	1.4453125	24.0207911	34.7175497
100G	15%	1	2.625	24.0207911	63.0545768
100G	27%	1	2.890625	24.0207911	69.4350994
150G	15%	1.3203125	3.9375	24.0207911	71.6359689
150G	27%	1.453125	4.3359375	24.0207911	71.6749413
200G	15%	1.7578125	5.25	24.0207911	71.7420962
200G	27%	2	4.40625	31.51	69.43
250G	15%	2.1953125	6	26.2727403	71.8059237
250G	27%	2.4140625	6	28.9312914	71.9068991
300G	15%	2.6328125	6	31.5272884	71.8485385
300G	27%	2.8984375	6	34.7175497	71.8681352
350G	15%	3.0703125	6	36.7818364	71.8790086

Trunk Payload	FEC	Min BPS	Max BPS	Min GBd	Max GBd
350G	27%	3.3828125	6	40.503808	71.8404724
400G	15%	3.5078125	6	42.0363845	71.9018782
400G	27%	3.8671875	6	46.2900663	71.8197392
450G	15%	3.9453125	6	47.2909326	71.9196757
450G	27%	4.34375	6	52.0763245	71.9327648
500G	15%	4.3828125	6	52.5454806	71.93392
500G	27%	4.8281250	6	57.8625828	71.9068991
550G	15%	4.8203125	6	57.8000287	71.9455787
550G	27%	5.3125	6	63.6488411	71.88575
600G	15%	5.2578125	-	-	71.9552971

The following table displays the trunk parameter ranges for the 1.2TL card.

Trunk Payload	FEC	Min BPS	Max BPS	Min GBd	Max GBd
100G	15%	1	2.625	24.0207911	63.0545768
100G	27%	1	2.890625	24.0207911	69.4350994
150G	15%	1.3203125	3.9375	24.0207911	71.6359689
150G	27%	1.453125	4.3359375	24.0207911	71.6749413
200G	15%	2	4	31.5272884	63.0545768
200G	27%	2	4.40625	31.51664088	69.43509943
250G	15%	2.1953125	4.5	35.0303204	71.8059237
250G	27%	2.4140625	4.5	38.5750552	71.9068991
300G	15%	2.6328125	4.5	42.0363845	71.8485385
300G	27%	2.8984375	4.5	46.2900662857142	71.86813526
350G	15%	3.0703125	4.5	49.0424486	71.8790086
350G	27%	3.3828125	4.5	54.0050773	71.8404724
400G	15%	3.5078125	4.5	56.0485127	71.9018782
400G	27%	3.8671875	4.5	61.72008838	71.81973921

To configure the BPS, see [Configuring the BPS, on page 23](#).

Configuring the Card Mode

You can configure the line card in the module (muxponder) or slice configuration (muxponder slice).

To configure the card in the muxponder mode, use the following commands.

configure

hw-module location *location* mxponder client-rate {100GE | OTU4}

hw-module location *location* mxponder trunk-rate {50G | 100G|150G | 200G | 250G | 300G | 350G | 400G | 450G | 500G | 550G | 600G }

commit

To configure the card in the muxponder slice mode, use the following commands.

configure

hw-module location *location* mxponder-slice *mxponder-slice-number* client-rate { 100GE|OTU4}

hw-module location *location* mxponder-slice trunk-rate { 100G | 200G | 300G | 400G | 500G | 600G }

commit

Examples

The following is a sample in which the card is configured in the muxponder mode with a 550G trunk payload.

```
RP/0/RP0/CPU0:ios#config
Tue Oct 15 01:24:56.355 UTC
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder client-rate 100GE
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder trunk-rate 550G
RP/0/RP0/CPU0:ios(config)#commit
```

The following is a sample in which the card is configured in the muxponder mode with a 500G trunk payload.

```
RP/0/RP0/CPU0:ios#config
Sun Feb 24 14:09:33.989 UTC
RP/0/RP0/CPU0:ios(config)#hw-module location 0/2 mxponder client-rate OTU4
RP/0/RP0/CPU0:ios(config)#hw-module location 0/2 mxponder trunk-rate 500G
RP/0/RP0/CPU0:ios(config)#commit
```

The following is a sample in which the card is configured in the muxponder slice 0 mode with a 500G trunk payload.

```
RP/0/RP0/CPU0:ios#config
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 0 client-rate 100GE
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 0 trunk-rate 500G
RP/0/RP0/CPU0:ios(config)#commit
```

The following is a sample in which the card is configured in the muxponder slice 1 mode with a 400G trunk payload.

```
RP/0/RP0/CPU0:ios#config
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 1 client-rate 100GE
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 1 trunk-rate 400G
RP/0/RP0/CPU0:ios(config)#commit
```

The following is a sample in which the card is configured with mixed client rates in the muxponder slice mode.

```
RP/0/RP0/CPU0:ios#configure
Mon Mar 23 06:10:22.227 UTC
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 0 client-rate OTU4 trunk-rate
500G
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 1 client-rate 100GE trunk-rate
500G
RP/0/RP0/CPU0:ios(config)#commit
```

Verifying the Card Configuration

```
RP/0/RP0/CPU0:ios#show hw-module location 0/2 mxponder
Fri Mar 15 11:48:48.344 IST
```

```
Location:                0/2
Client Bitrate:          100GE
Trunk Bitrate:           500G
Status:                  Provisioned
LLDP Drop Enabled:       FALSE
Client Port              Mapper/Trunk Port   CoherentDSP0/2/0/0   CoherentDSP0/2/0/1
                        Traffic Split Percentage

HundredGigEctrler0/2/0/2   ODU40/2/0/0/1           100                   0
HundredGigEctrler0/2/0/3   ODU40/2/0/0/2           100                   0
HundredGigEctrler0/2/0/4   ODU40/2/0/0/3           100                   0
HundredGigEctrler0/2/0/5   ODU40/2/0/0/4           100                   0
HundredGigEctrler0/2/0/6   ODU40/2/0/0/5           100                   0
HundredGigEctrler0/2/0/7   ODU40/2/0/1/1           0                     100
HundredGigEctrler0/2/0/8   ODU40/2/0/1/2           0                     100
HundredGigEctrler0/2/0/9   ODU40/2/0/1/3           0                     100
HundredGigEctrler0/2/0/10  ODU40/2/0/1/4           0                     100
HundredGigEctrler0/2/0/11  ODU40/2/0/1/5           0                     100
```

The following is a sample output of the coupled mode configuration where the shared client port is highlighted.

```
RP/0/RP0/CPU0:ios#show hw-module location 0/1 mxponder
Tue Oct 15 01:25:57.358 UTC
```

```
Location:                0/1
Client Bitrate:          100GE
Trunk Bitrate:           550G
Status:                  Provisioned
LLDP Drop Enabled:       FALSE
Client Port              Mapper/Trunk Port   CoherentDSP0/1/0/0   CoherentDSP0/1/0/1
                        Traffic Split Percentage

HundredGigEctrler0/1/0/2   ODU40/1/0/0/1           100                   0
HundredGigEctrler0/1/0/3   ODU40/1/0/0/2           100                   0
HundredGigEctrler0/1/0/4   ODU40/1/0/0/3           100                   0
HundredGigEctrler0/1/0/5   ODU40/1/0/0/4           100                   0
HundredGigEctrler0/1/0/6   ODU40/1/0/0/5           100                   0
HundredGigEctrler0/1/0/7   ODU40/1/0/0/6           50                    50
HundredGigEctrler0/1/0/8   ODU40/1/0/1/1           0                     100
HundredGigEctrler0/1/0/9   ODU40/1/0/1/2           0                     100
HundredGigEctrler0/1/0/10  ODU40/1/0/1/3           0                     100
HundredGigEctrler0/1/0/11  ODU40/1/0/1/4           0                     100
HundredGigEctrler0/1/0/12  ODU40/1/0/1/5           0                     100
```

The following is a sample output of all the muxponder slice 0 configurations.

```
RP/0/RP0/CPU0:ios#show hw-module location 0/1 mxponder-slice 0
Fri Mar 15 06:04:18.348 UTC
```

```
Location:                0/1
```

```

Slice ID:                0
Client Bitrate:          100GE
Trunk Bitrate:           500G
Status:                  Provisioned
LLDP Drop Enabled:      FALSE
Client Port              Mapper/Trunk Port      CoherentDSP0/1/0/0
                        Traffic Split Percentage

HundredGigECtrlr0/1/0/2    ODU40/1/0/0/1          100
HundredGigECtrlr0/1/0/3    ODU40/1/0/0/2          100
HundredGigECtrlr0/1/0/4    ODU40/1/0/0/3          100
HundredGigECtrlr0/1/0/5    ODU40/1/0/0/4          100
HundredGigECtrlr0/1/0/6    ODU40/1/0/0/5          100

```

The following is a sample output of all the muxponder slice 1 configurations.

```

RP/0/RP0/CPU0:ios#show hw-module location 0/1 mxponder-slice 1
Fri Mar 15 06:11:50.020 UTC

Location:                0/1
Slice ID:                 1
Client Bitrate:           100GE
Trunk Bitrate:            400G
Status:                   Provisioned
LLDP Drop Enabled:        TRUE
Client Port              Mapper/Trunk Port      CoherentDSP0/1/0/1
                        Traffic Split Percentage

HundredGigECtrlr0/1/0/8    ODU40/1/0/1/1          100
HundredGigECtrlr0/1/0/9    ODU40/1/0/1/2          100
HundredGigECtrlr0/1/0/10   ODU40/1/0/1/3          100
HundredGigECtrlr0/1/0/11   ODU40/1/0/1/4          100

```

The following is a sample output of the muxponder slice 1 configuration with client configured as OTU4.

```

RP/0/RP0/CPU0:ios#sh hw-module location 0/0 mxponder-slice 1
Wed Mar 11 13:59:11.073 UTC

Location:                0/0
Slice ID:                 1
Client Bitrate:           OTU4
Trunk Bitrate:            200G
Status:                   Provisioned
Client Port              Peer/Trunk Port      CoherentDSP0/0/0/1
                        Traffic Split Percentage

OTU40/0/0/8               ODU40/0/0/1/1          100
OTU40/0/0/9               ODU40/0/0/1/2          100

```

The following is a sample to verify the mixed client rate configuration in the muxponder slice mode.

```

RP/0/RP0/CPU0:ios#show hw-module location 0/1 mxponder
Mon Mar 23 06:20:22.227 UTC

Location:                0/1
Slice ID:                 0
Client Bitrate:           OTU4
Trunk Bitrate:            500G
Status:                   Provisioned
Client Port              Peer/Trunk Port      CoherentDSP0/1/0/0
                        Traffic Split Percentage

OTU40/1/0/2               ODU40/1/0/0/1          100
OTU40/1/0/3               ODU40/1/0/0/2          100

```

```

OTU40/1/0/4                ODU40/1/0/0/3                100
OTU40/1/0/5                ODU40/1/0/0/4                100
OTU40/1/0/6                ODU40/1/0/0/5                100

Location:                   0/1
Slice ID:                    1
Client Bitrate:              100GE
Trunk Bitrate:               500G
Status:                       Provisioned
LLDP Drop Enabled:           FALSE
ARP Snoop Enabled:           FALSE
Client Port                   Mapper/Trunk Port              CoherentDSP0/1/0/1
                              Traffic Split Percentage

HundredGigECtrlr0/1/0/8    ODU40/1/0/1/1                100
HundredGigECtrlr0/1/0/9    ODU40/1/0/1/2                100
HundredGigECtrlr0/1/0/10   ODU40/1/0/1/3                100
HundredGigECtrlr0/1/0/11   ODU40/1/0/1/4                100
HundredGigECtrlr0/1/0/12   ODU40/1/0/1/5                100

```

Use the following command to clear alarm statistics on the optics or coherent DSP controller.

clear counters controller *controllertype* R/S/I/P

The following is a sample in which the alarm statistics are cleared on the coherent DSP controller.

```

RP/0/RP0/CPU0:ios#show controller coherentDSP 0/1/0/0
Tue Jun 11 05:15:12.540 UTC

Port                               : CoherentDSP 0/1/0/0
Controller State                     : Up
Inherited Secondary State             : Normal
Configured Secondary State           : Normal
Derived State                         : In Service
Loopback mode                         : None
BER Thresholds                       : SF = 1.0E-5  SD = 1.0E-7
Performance Monitoring                : Enable

Alarm Information:
LOS = 1 LOF = 1 LOM = 0
OOF = 1 OOM = 1 AIS = 0
IAE = 0 BIAE = 0          SF_BER = 0
SD_BER = 2      BDI = 2 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0
Detected Alarms                       : None

Bit Error Rate Information
PREFEC BER                             : 8.8E-03
POSTFEC BER                             : 0.0E+00

TTI :
  Remote hostname                       : P2B8
  Remote interface                       : CoherentDSP 0/1/0/0
  Remote IP addr                         : 0.0.0.0

FEC mode                               : Soft-Decision 15

AINS Soak                              : None
AINS Timer                              : 0h, 0m
AINS remaining time                     : 0 seconds
RP/0/RP0/CPU0:ios#clear counters controller coherentDSP 0/1/0/0
Tue Jun 11 05:17:07.271 UTC
All counters are cleared

```

```

RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/1/0/1
Tue Jun 11 05:20:55.199 UTC

Port                               : CoherentDSP 0/1/0/1
Controller State                    : Up
Inherited Secondary State          : Normal
Configured Secondary State        : Normal
Derived State                      : In Service
Loopback mode                      : None
BER Thresholds                     : SF = 1.0E-5  SD = 1.0E-7
Performance Monitoring             : Enable

Alarm Information:
LOS = 0 LOF = 0 LOM = 0
OCF = 0 OOM = 0 AIS = 0
IAE = 0 BIAE = 0          SF_BER = 0
SD_BER = 0      BDI = 0 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0
Detected Alarms                  : None

Bit Error Rate Information
PREFEC BER                       : 1.2E-02
POSTFEC BER                      : 0.0E+00

TTI :
  Remote hostname                 : P2B8
  Remote interface                : CoherentDSP 0/1/0/1
  Remote IP addr                  : 0.0.0.0

FEC mode                          : Soft-Decision 15

AINS Soak                        : None
AINS Timer                       : 0h, 0m
AINS remaining time              : 0 seconds

```

Regeneration Mode

In an optical transmission system, 3R regeneration helps extend the reach of the optical communication links by reamplifying, reshaping, and retiming the data pulses. Regeneration helps to correct any distortion of optical signals by converting it to an electrical signal, processing that electrical signal, and then retransmitting it again as an optical signal.

In Regeneration (Regen) mode, the OTN signal is received on a trunk port and the regenerated OTN signal is sent on the other trunk port of the line card and the other way round. In this mode, only the trunk optics controller and coherentDSP controllers are created.

Configuring the Card in Regen Mode

The supported trunk rates for the different cards are:

To configure regen mode on 1.2T, 1.2TL, and 2-QDD-C cards, use the following commands:

```

configure
hw-module location location
regen
trunk-rate trunk-rate
commit

```

exit

Example

Verifying the Regen Mode

The following is a sample to verify the regen mode.

show hw-module location *location* regen

```
RP/0/RP0/CPU0:ios#show hw-module location 0/0 regen
Mon Mar 25 09:50:42.936 UTC

Location:                0/0
Trunk Bitrate:           400G
Status:                  Provisioned
East Port                 West Port
CoherentDSP0/0/0/0      CoherentDSP0/0/0/1
```

The terms, East Port and West Port are used to represent OTN signal regeneration at the same layer.

Configuring the BPS

You can configure the Bits per Symbol (BPS) to 3.4375 to support 300G trunk configurations on 75 GHz networks using the following commands:

configure

controller optics *R/S/I/P* bits-per-symbol 3.4375

commit

The following is a sample in which the BPS is configured to 3.4375.

```
RP/0/RP0/CPU0:ios#configure
Wed Mar 27 14:12:49.932 UTC
RP/0/RP0/CPU0:ios(config)#controller optics 0/3/0/0 bits-per-symbol 3.4375
RP/0/RP0/CPU0:ios(config)#commit
```

Viewing BPS and Baud Rate Ranges

To view the the BPS for a specific range use the following command:

show controller optics *R/S/I/P* bps-range *bps-range* | include *data-rate* | include *fec-type*

```
RP/0/RP0/CPU0:ios#show controllers optics 0/3/0/0 bps-range 3 3.05 | include 300G | include
SD27
Thu Mar 28 03:01:39.751 UTC
300G          SD27          3.0000000          69.4350994
300G          SD27          3.0078125          69.2547485
300G          SD27          3.0156250          69.0753320
300G          SD27          3.0234375          68.8968428
300G          SD27          3.0312500          68.7192736
300G          SD27          3.0390625          68.5426174
300G          SD27          3.0468750          68.3668671
```

To view the baud for a specific range use the following command:

show controller optics *R/S/I/P* baud-rate-range *baud-range* | include *data-rate* | include *fec-type*

```
RP/0/RP0/CPU0:ios#show controllers optics 0/3/0/0 baud-rate-range 43 43.4 | include 300G |
include SD27
Thu Mar 28 03:12:36.521 UTC
300G          SD27          4.8046875      43.3545986
300G          SD27          4.8125000      43.2842178
300G          SD27          4.8203125      43.2140651
300G          SD27          4.8281250      43.1441394
300G          SD27          4.8359375      43.0744397
300G          SD27          4.8437500      43.0049648
```

Configuring the Trunk Rate for BPSK

From R7.2.1 onwards, you can configure trunk rates of 50G, 100G, and 150G to support Binary Phase-Shift Keying (BPSK) modulation. The BPSK modulation enables information to be carried over radio signals more efficiently.

You can configure trunk rates for BPSK using CLI, NetConf YANG, and OC models.

The following table lists the 50G, 100G, and 150G trunk rates with the supported BPSK modulation:

Trunk Rate	BPSK Modulation
50G	1 to 1.4453125
100G	1 to 2.890625
150G	1.453125 to 4.3359375

To configure the trunk rate for BPSK modulation, enter the following commands:

```
configure
hw-module location location mxponder
trunk-rate {50G | 100G | 150G}
commit
```

The following example shows how to configure trunk rate to 50G:

```
RP/0/RP0/CPU0:(config)#hw-module location 0/0 mxponder
RP/0/RP0/CPU0:(config-hwmod-mxp)#trunk-rate 50G
RP/0/RP0/CPU0:(config-hwmod-mxp)#commit
```

Viewing the BPSK Trunk Rate Ranges

To view the trunk rate configured for the BPSK modulation, use the following **show** commands:

```
RP/0/RP0/CPU0:ios (hwmod-mxp) #show hw-module location 0/0 mxponder

Tue Feb 25 11:13:41.934 UTC

Location:          0/0
Client Bitrate:    100GE
Trunk Bitrate:     50G
Status:            Provisioned
LLDP Drop Enabled: FALSE
ARP Snoop Enabled: FALSE
```

```

Client Port                               Mapper/Trunk Port       CoherentDSP0/0/0/0
CoherentDSP0/0/0/1                       Traffic Split Percentage

HundredGigECtrlr0/0/0/2                 ODU40/0/0/0           50
50
    
```

```

RP/0/RP0/CPU0:ios#show controllers optics 0/0/0/0
Thu Mar  5 07:12:55.681 UTC
    
```

```

Controller State: Up
Transport Admin State: In Service
Laser State: On
LED State: Green
    
```

Optics Status

```

Optics Type: DWDM optics
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 = 2
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 %
Actual TX Power = 1.97 dBm
RX Power = 1.58 dBm
RX Signal Power = 0.60 dBm
Frequency Offset = 386 MHz
    
```

Performance Monitoring: Enable

THRESHOLD VALUES

Parameter	High Alarm	Low Alarm	High Warning	Low Warning
Rx Power Threshold(dBm)	4.9	-12.0	0.0	0.0
Tx Power Threshold(dBm)	3.5	-10.1	0.0	0.0
LBC Threshold(mA)	N/A	N/A	0.00	0.00

```

Configured Tx Power = 2.00 dBm
Configured CD High Threshold = 180000 ps/nm
Configured CD lower Threshold = -180000 ps/nm
Configured OSNR lower Threshold = 0.00 dB
    
```

```

Configured DGD Higher Threshold = 180.00 ps
Baud Rate = 34.7175521851 GBd
Bits per Symbol = 1.0000000000 bits/symbol
Modulation Type: BPSK
Chromatic Dispersion -9 ps/nm
Configured CD-MIN -180000 ps/nm CD-MAX 180000 ps/nm
Polarization Mode Dispersion = 0.0 ps
Second Order Polarization Mode Dispersion = 125.00 ps^2
Optical Signal to Noise Ratio = 34.60 dB
SNR = 20.30 dB
Polarization Dependent Loss = 0.20 dB
Polarization Change Rate = 0.00 rad/s
Differential Group Delay = 2.00 ps
Filter Roll Off Factor : 0.100
Rx VOA Fixed Ratio : 15.00 dB
Enhanced Colorless Mode : 0
Enhanced SOP Tolerance Mode : 0
NLEQ Compensation Mode : 0
Cross Polarization Gain Mode : 0
Cross Polarization Weight Mode : 0
Carrier Phase Recovery Window : 0
Carrier Phase Recovery Extended Window : 0
    
```

```

AINS Soak           : None
AINS Timer          : 0h, 0m
AINS remaining time : 0 seconds
    
```

OTN-XP Card

The following section describes the supported configurations and procedures to configure the card modes on the line card.

LC Mode on OTN-XP Card

When you install the OTN-XP card in the Cisco NCS 1004 chassis, it is in the POWERED_ON state. The **LCMODE is not configured for line card** alarm is present on the card and the LED status is AMBER.

```

sysadmin-vm:0_RP0# show platform
Thu Mar 26 21:38:07.305 UTC+00:00
Location  Card Type                HW State    SW State    Config State
-----
0/0      NCS1K4-LC-FILLER         PRESENT     N/A         NSHUT
0/1      NCS1K4-OTN-XP           POWERED_ON  N/A         NSHUT
0/RP0    NCS1K4-CNTRLR-K9        OPERATIONAL OPERATIONAL  NSHUT
0/FT0    NCS1K4-FAN              OPERATIONAL N/A         NSHUT
0/FT1    NCS1K4-FAN              OPERATIONAL N/A         NSHUT
0/FT2    NCS1K4-FAN              OPERATIONAL N/A         NSHUT
0/PM0    NCS1K4-AC-PSU           OPERATIONAL N/A         NSHUT
0/SC0    NCS1004                  OPERATIONAL N/A         NSHUT

sysadmin-vm:0_RP0# show alarms brief system active
Thu Mar 26 21:38:34.394 UTC+00:00

-----
Active Alarms
-----
Location          Severity          Group            Set time          Description
-----
    
```

```

0          major          environ          03/26/20 20:23:11  Power Module redundancy
lost.
0          critical       environ          03/26/20 20:23:29  Fan: One or more LCs
missing, running fans at max speed.
0/1       not_alarmed    shelf          03/26/20 21:38:26  LCMODE is not configured
for line card
sysadmin-vm:0_RP0#

sysadmin-vm:0_RP0# show led location 0/1
Thu Mar 26 21:39:05.101 UTC+00:00
=====
Location  LED Name                               Mode          Color
=====
0/1
          0/1-Status LED                          WORKING       AMBER
sysadmin-vm:0_RP0#

```

You must select a datapath mode by configuring the LC mode. OTN-XP does not have a default LC mode. After the LC mode is configured using the CLI, the card transitions to the OPERATIONAL state, the alarm clears, and the LED status turns to GREEN.

The LC modes supported on the OTN-XP card are:

- 10G-GREY-MXP



Note 100G-TXP LC mode is not supported.

Only one LC mode can be configured on the OTN-XP card at a time. When the LC mode is changed using the CLI, the **LCMODE changed, delete the datapath config and reload line card** alarm is present on the card and the DP FPD is in disabled state. To clear the alarm and enable the DP FPD, delete the existing datapath configuration and reload the line card to apply the new LC mode to make the card operational.

If a LC mode requires a different FPGA configuration, and the package is not available, the **OTN_XP_DP_FPD_PKG is missing, please install the package to proceed** alarm is present on the card. To clear the alarm, install the OTN_XP_DP_FPD_PKG file. After the package installation is complete, the required FPGA image is copied from the OTN_XP_DP_FPD_PKG file to the card, the card is automatically reloaded, and the card becomes operational.



Note The LC mode configuration is a shared plane configuration. The configuration does not enter the preconfigured state when the line card is not available.

Configuring the LC Mode



Note

- Ensure the OTN_XP_DP_FPD_PKG file is installed before configuring the LC mode.
- When you insert an OTN-XP line card having a lower FPD version, you must configure a LC mode which is supported on the software release that the line card is loaded with. You cannot upgrade the FPD of a line card if you configure a LC mode supported only in a higher software release.

To configure the LC mode on the OTN-XP card, use the following commands:

configure**lc-module location** *location* **lcmode** *mode***commit****Example**

To view the LC modes available on the OTN-XP card, use the following command:

```
RP/0/RP0/CPU0:ios#sh lc-module location 0/0 lcmode all
Wed Sep 29 14:41:51.487 UTC
States: A-Available      R-Running      C-Configured
```

Node	Lcmode_Supported	Owner	Options(State)	HW_Ver
0/0	Yes	None	10G-GREY-MXP (A) 4x100G-MXP-400G-TXP (A)	3.0 2.0

The following is a sample in which the OTN-XP card is configured in the 10G-GREY-MXP mode.

```
RP/0/RP0/CPU0:ios#configure
Thu Mar 26 21:40:51.495 UTC
RP/0/RP0/CPU0:ios(config)#lc-module location 0/1 lcmode 10G-GREY-MXP
RP/0/RP0/CPU0:ios(config)#commit
```

Verifying the LC Mode Configuration

The following is a sample output of a successful 10G-GREY-MXP LC mode configuration after which the card transitions to the OPERATIONAL state, the alarm clears, and the LED status turns to GREEN.

```
RP/0/RP0/CPU0:ios(config)#do show platform
Thu Mar 26 21:41:17.206 UTC
```

Node	Type	State	Config state
0/0	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/1	NCS1K4-OTN-XP	OPERATIONAL	NSHUT
0/RP0/CPU0	NCS1K4-CNTRLR-K9(Active)	IOS XR RUN	NSHUT
0/FT0	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT1	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT2	NCS1K4-FAN	OPERATIONAL	NSHUT
0/PM0	NCS1K4-AC-PSU	OPERATIONAL	NSHUT
0/SC0	NCS1004	OPERATIONAL	NSHUT

```
RP/0/RP0/CPU0:ios(config)#do show alarms brief system active
Thu Mar 26 21:41:29.641 UTC
```

Active Alarms

Location	Severity	Group	Set Time	Description
0	Major	Environ	03/26/2020 20:23:11 UTC	Power Module redundancy lost.
0	Critical	Environ	03/26/2020 20:23:29 UTC	Fan: One or more LCs missing, running fans at max speed.

```
RP/0/RP0/CPU0:ios(config)#end
RP/0/RP0/CPU0:ios#show lc-module location 0/1 lcmode all
Thu Mar 26 21:41:58.780 UTC
```

```

States: A-Available      R-Running      C-Configured

Node  Lcmode_Supported  Owner  Options(State)  HW_Ver
-----
0/1   Yes              CLI    10G-GREY-MXP (R/C)  3.0
              4x100G-MXP-400G-TXP (A)  2.0
RP/0/RP0/CPU0:ios#show lc-module location 0/1 lcmode
Thu Mar 26 21:42:18.997 UTC

Node  Lcmode_Supported  Owner  Running  Configured
-----
0/1   Yes              CLI    10G-GREY-MXP      10G-GREY-MXP
RP/0/RP0/CPU0:ios#admin
Thu Mar 26 21:42:38.525 UTC

root connected from 192.0.2.3 using ssh on sysadmin-vm:0_RP0
sysadmin-vm:0_RP0# show led location 0/1
Thu Mar 26 21:42:45.337 UTC+00:00
=====
Location  LED Name  Mode  Color
=====
0/1
          0/1-Status LED      WORKING  GREEN
    
```

Example

The following is a sample in which the LC mode is changed from 10G-GREY-MXP to the 4x100G-MXP-400G-TXP mode. In this sample, the datapath configuration is deleted and the card is reloaded to apply the new LC mode.

```

RP/0/RP0/CPU0:ios#show lc-module location all lcmode
Thu Sep 30 10:19:29.853 UTC

Node  Lcmode_Supported  Owner  Running  Configured
-----
0/0   Yes              CLI    10G-GREY-MXP      10G-GREY-MXP
0/1   No                N/A    N/A        N/A
0/2   No                N/A    N/A        N/A
0/3   No                N/A    N/A        N/A

RP/0/RP0/CPU0:ios#configure
Thu Sep 30 10:19:32.818 UTC
Current Configuration Session  Line      User      Date              Lock
00001000-000051f7-00000000   vty1     root      Wed Sep 29 15:26:00 2021
RP/0/RP0/CPU0:ios(config)#no lc-module location 0/0 lcmode 10g-GREY-MXP
RP/0/RP0/CPU0:ios(config)#commit
Thu Sep 30 10:20:34.086 UTC
RP/0/RP0/CPU0:ios(config)#do show alarms brief system active
Thu Sep 30 10:20:52.950 UTC

-----
Active Alarms
-----
Location      Severity  Group      Set Time          Description
-----
0/PM0        Major     Environ    09/29/2021 14:41:59 UTC  Power Module Output
    
```

Disabled

0 Major Environ 09/29/2021 14:42:15 UTC Power Module
redundancy lost.

0 Critical Environ 09/29/2021 14:42:25 UTC Fan: One or more
LCs missing, running fans at max speed.

0/0 NotAlarmed Shelf 09/30/2021 10:20:34 UTC LCMODE changed,
delete the datapath config and reload line card

```
RP/0/RP0/CPU0:ios#configure
Thu Sep 30 10:21:41.281 UTC
Current Configuration Session Line User Date Lock
00001000-000051f7-00000000 vty1 root Wed Sep 29 15:26:00 2021
RP/0/RP0/CPU0:ios(config)#no hw-module location 0/0
RP/0/RP0/CPU0:ios(config)#commit
Thu Sep 30 10:21:49.982 UTC
RP/0/RP0/CPU0:ios(config)#
```

```
RP/0/RP0/CPU0:ios#show platform
Thu Sep 30 10:22:08.482 UTC
```

Node	Type	State	Config state
------	------	-------	--------------

0/0	NCS1K4-OTN-XP	OPERATIONAL	NSHUT
0/2	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/3	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/RP0/CPU0	NCS1K4-CNTRLR-K9(Active)	IOS XR RUN	NSHUT
0/FT0	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT1	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT2	NCS1K4-FAN	OPERATIONAL	NSHUT
0/PM0	NCS1K4-AC-PSU	OPERATIONAL	NSHUT
0/SC0	NCS1004	OPERATIONAL	NSHUT

```
RP/0/RP0/CPU0:ios#
```

```
RP/0/RP0/CPU0:ios#admin
```

```
Thu Sep 30 10:23:55.937 UTC
```

```
Last login: Thu Sep 30 04:32:57 2021 from 192.0.2.3
```

```
root connected from 192.0.2.3 using ssh on sysadmin-vm:0_RP0
```

```
sysadmin-vm:0_RP0# hw-module location 0/0 reload
```

```
Thu Sep 30 10:24:17.938 UTC+00:00
```

```
Reloading the module will be traffic impacting if not properly drained. Continue to Reload  
hardware module ? [no,yes] yes
```

```
result Card graceful reload request on 0/0 succeeded.
```

```
sysadmin-vm:0_RP0#show platform
```

```
Thu Sep 30 10:25:16.876 UTC+00:00
```

Location	Card Type	HW State	SW State	Config State
----------	-----------	----------	----------	--------------

0/0	NCS1K4-OTN-XP	POWERED_ON	N/A	NSHUT
0/2	NCS1K4-LC-FILLER	PRESENT	N/A	NSHUT
0/3	NCS1K4-LC-FILLER	PRESENT	N/A	NSHUT
0/RP0	NCS1K4-CNTRLR-K9	OPERATIONAL	OPERATIONAL	NSHUT
0/FT0	NCS1K4-FAN	OPERATIONAL	N/A	NSHUT
0/FT1	NCS1K4-FAN	OPERATIONAL	N/A	NSHUT
0/FT2	NCS1K4-FAN	OPERATIONAL	N/A	NSHUT
0/PM0	NCS1K4-2KW-AC	OPERATIONAL	N/A	NSHUT
0/SC0	NCS1004-K9	OPERATIONAL	N/A	NSHUT

```
sysadmin-vm:0_RP0#exit
```

```
RP/0/RP0/CPU0:ios#show lc-module location all lcmode
Thu Sep 30 10:29:08.183 UTC
```

Node	Lcmode_Supported	Owner	Running	Configured
0/0	Yes	None	Not running	Not configured
0/1	No	N/A	N/A	N/A
0/2	No	N/A	N/A	N/A
0/3	No	N/A	N/A	N/A

```
RP/0/RP0/CPU0:ios#show platform
Thu Sep 30 10:29:36.075 UTC
```

Node	Type	State	Config state
0/0	NCS1K4-OTN-XP	POWERED_ON	NSHUT
0/2	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/3	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/RP0/CPU0	NCS1K4-CNTLR-K9 (Active)	IOS XR RUN	NSHUT
0/FT0	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT1	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT2	NCS1K4-FAN	OPERATIONAL	NSHUT
0/PM0	NCS1K4-AC-PSU	OPERATIONAL	NSHUT
0/SC0	NCS1004	OPERATIONAL	NSHUT

```
RP/0/RP0/CPU0:ios#
```

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

```
Thu Sep 30 10:29:57.997 UTC
```

```
Current Configuration Session Line User Date Lock
```

```
00001000-000051f7-00000000 vty1 root Wed Sep 29 15:26:00 2021
```

```
RP/0/RP0/CPU0:ios(config)#lc-module location 0/0 lcmode 4x100G-MXP-400G-TXP
```

```
RP/0/RP0/CPU0:ios(config)#commit
```

```
Thu Sep 30 10:30:11.312 UTC
```

```
RP/0/RP0/CPU0:ios(config)#end
```

```
RP/0/RP0/CPU0:ios#show lc-module location all lcmode
```

```
Thu Sep 30 10:40:56.480 UTC
```

Node	Lcmode_Supported	Owner	Running	Configured
0/0	Yes	CLI	4x100G-MXP-400G-TXP	4x100G-MXP-400G-TXP
0/1	No	N/A	N/A	N/A
0/2	No	N/A	N/A	N/A
0/3	No	N/A	N/A	N/A

```
RP/0/RP0/CPU0:ios# RP/0/RP0/CPU0:ios#show platform
```

```
Thu Sep 30 10:41:25.093 UTC
```

Node	Type	State	Config state
0/0	NCS1K4-OTN-XP	OPERATIONAL	NSHUT
0/2	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/3	NCS1K4-LC-FILLER	PRESENT	NSHUT
0/RP0/CPU0	NCS1K4-CNTLR-K9 (Active)	IOS XR RUN	NSHUT
0/FT0	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT1	NCS1K4-FAN	OPERATIONAL	NSHUT
0/FT2	NCS1K4-FAN	OPERATIONAL	NSHUT
0/PM0	NCS1K4-AC-PSU	OPERATIONAL	NSHUT
0/SC0	NCS1004	OPERATIONAL	NSHUT

```
RP/0/RP0/CPU0:ios#
```

Muxponder Configuration on OTN-XP Card

The OTN-XP card has two trunk ports and 12 client ports. The muxponder configuration supports two slices, 0 and 1. You can configure mxponder-slice 0, mxponder-slice 1, or both. Each mxponder-slice supports 10 client interfaces.

Table 3: Feature History

Feature Name	Release Information	Description
400 TXP or MXP modes with CFP2 DCO for OTN-XP Card	Cisco IOS XR Release 7.3.1	<p>On the OTN-XP card, you can configure a single 400GE or 4x100G payload that is received over the client port as a 400G signal over DWDM on the line side.</p> <p>The card improves efficiency, performance, and flexibility for customer networks allowing 400GE or 4x100G client transport over 400G WDM wavelength.</p> <p>Commands modified:</p> <ul style="list-style-type: none"> • controller coherentDSP • show controller coherentDSP

Table 4: Hardware Module Configuration with Client to Trunk Mapping

Hardware Module Configuration	Line Card Mode	Client Port Rate	Client to Trunk Mapping	Trunk Rate
10G Grey Muxponder	10G-GREY-MXP	OTU2, OTU2e, or 10 GE	<p>Mxponder-slice 0—Client ports 4, 5, and 2 are mapped to the trunk port 0.</p> <p>Mxponder-slice 1—Client ports 7, 6, and 11 are mapped to the trunk port 1.</p> <p>Each client port consists of four lanes, 1, 2, 3, and 4. The lanes 3 and 4 can only be configured for ports 2 and 11. It is not mandatory to configure all 10 client lanes for a slice.</p>	100G

Configuring the Muxponder Mode for 10G Grey Muxponder



Note The LC mode must be configured to 10G-GREY-MXP on the OTN-XP card before you perform this configuration.

To configure the OTN-XP card in the muxponder mode, use the following commands:

configure

hw-module location *location* **mxponder-slice** *mxponder-slice-number*

trunk-rate 100G

client-port-rate *client-port-number* **lane** *lane-number* **client-type** { 10GE | OTU2 | OTU2e }

commit

Example

The following is a sample in which the OTN-XP card is configured with mixed client rates in the mxponder-slice 0 mode.

```
RP/0/RP0/CPU0:ios#config
Tue Apr 21 09:21:44.460 UTC
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 0
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#trunk-rate 100G
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#client-port-rate 2 lane 3 client-type OTU2
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#client-port-rate 2 lane 4 client-type OTU2
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#client-port-rate 4 lane 1 client-type 10GE
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#commit
```

Verifying the Muxponder Configuration

The following is a sample to verify the muxponder configuration in the OTN-XP card.

```
RP/0/RP0/CPU0:ios#show hw-module location 0/1 mxponder
Tue Apr 21 09:26:12.308 UTC

Location:                0/1
Slice ID:                 0
Client Bitrate:          MIXED
Trunk Bitrate:           100G
Status:                  Provisioned
LLDP Drop Enabled:      FALSE
ARP Snoop Enabled:      FALSE
Client Port Mapper/Trunk Port Peer/Trunk Port OTU40/0/0/0
Traffic Split Percentage
OTU20/0/0/0/2/3          NONE ODU20/0/0/0/2/3          100
OTU20/0/0/0/2/4          NONE ODU20/0/0/0/2/4          100
TenGigECtrlr0/0/0/4/1 ODU2E0/0/0/0/4/1          NONE          100
```




CHAPTER 4

Configuring Controllers

There are three types of controllers for the line card. The controllers are the optics controller, the ethernet controller, and the coherent DSP controller. This chapter describes the procedures used to configure these controllers.



Note Unless otherwise specified, “line cards” refers to 1.2T and 1.2TL line cards.

- [AINS, on page 35](#)
- [FEC, on page 44](#)
- [Laser Squelching, on page 49](#)
- [Idle Insertion, on page 51](#)
- [LLDP Drop, on page 53](#)
- [Link Layer Discovery Protocol \(LLDP\) Support on Management Interface, on page 56](#)
- [MAC Address Snooping on Client Ports, on page 60](#)
- [Loopback, on page 62](#)
- [Restore Factory Settings, on page 68](#)
- [Headless Mode, on page 69](#)
- [Trail Trace Identifier, on page 69](#)
- [Chromatic dispersion, on page 71](#)
- [Frequency, on page 73](#)
- [Pseudo Random Binary Sequence, on page 73](#)

AINS

The Automatic-In-Service (AINS) feature allows the controller to automatically move to the automatic-in-service state after the maintenance window is completed. A soak time period is associated with the AINS state. The controller automatically moves to the In-Service state after the soak time period is completed. During the AINS maintenance window, alarms are not propagated to the EMS/NMS monitoring system.

You can configure AINS on the client ports of the card.

AINS States

The following table lists the AINS states.

State	Description
None	AINS is not enabled on the controller or the soak time period is complete.
Pending	AINS is configured on the controller. However, the soak time period has not started because either the primary state of controller is in Shutdown, Admin down, or Not ready state or the secondary state is in Maintenance state. AINS can also move to Pending state if alarms are raised during the soak time period.
Running	AINS is enabled on the controller. The primary state of the controller is Up and the secondary state is AINS.

If there are any service-affecting alarms when AINS is running on ethernet or optics controllers, the AINS state moves to Pending state. When the alarms are cleared, the AINS state moves to Running state.

The AINS soak time period restarts when there are line card reloads, XR reloads, line card warm reloads, power cycles, or alarm conditioning.

Soak Time Period

You can configure the soak time period to be between 1 minute to 48 hours.

All alarms are suppressed during the AINS state. When the optical and ethernet alarms are raised on the port during the soak time period, the AINS state moves to Pending. These alarms are not displayed in the output of the **show alarms brief card location 0/RP0/CPU0 active** command but in the output of the **show alarms brief card location 0/RP0/CPU0 conditions** command. When all the alarms clear, the soak time period starts, and the AINS state moves to Running. When the soak time period expires, the port moves to IS state.

Configuring AINS

To configure AINS on a muxponder, use the following command:

configure

hw-module location *location* **mxponder client-port-ains-soak** **hours** *hours* **minutes** *minutes*

commit

The following is a sample in which all client ports are configured with AINS with soak time period specified to be 15 minutes.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#hw-module location 0/3 mxponder client-port-ains-soak hours 0
minutes 15
RP/0/RP0/CPU0:ios(config)#commit
```

To configure AINS on a muxponder slice, use the following command:

configure

hw-module location *location* **mxponder-slice** *slice-number* **client-port-ains-soak** **hours** *hours* **minutes** *minutes*

commit

The following is a sample in which slice 0 client ports are configured with AINS with soak time period specified to be 40 minutes.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#hw-module location 0/3 mxponder-slice 0 client-port-ains-soak
hours 0 minutes 40
RP/0/RP0/CPU0:ios(config)#commit
```

Disabling AINS

To disable AINS on all muxponder client ports, set the hours and minutes to 0. Use the following commands:

configure

hw-module location *location* **mxponder** **client-port-ains-soak** **hours** *hours* **minutes** *minutes*

commit

The following is a sample in which AINS is disabled on all client ports.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#hw-module location 0/3 mxponder client-port-ains-soak hours 0
minutes 0
RP/0/RP0/CPU0:ios(config)#commit
```

To disable AINS on a muxponder slice, set the hours and minutes to 0. Use the following command:

configure

hw-module location *location* **mxponder-slice** *slice-number* **client-port-ains-soak** **hours** *hours* **minutes** *minutes*

commit

The following is a sample in which AINS is disabled on all client ports of slice 0.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#hw-module location 0/3 mxponder-slice 0 client-port-ains-soak
hours 0 minutes 0
RP/0/RP0/CPU0:ios(config)#commit
```

Displaying the AINS Configuration

The AINS Soak field in the output indicates the current state of AINS. The current state can be None, Pending, or Running. The Total Duration field indicates the total soak time period that is configured. The Remaining Duration field indicates the soak time that remains, after which, the AINS state moves to None.

This example displays the ethernet controller statistics with AINS Soak in running state.

```
RP/0/RP0/CPU0:ios#show controller HundredGigECtrlr 0/1/0/2
Thu Feb 21 19:52:55.001 UTC
Operational data for interface HundredGigECtrlr0/1/0/2:
State:
```

```

Administrative state: enabled
Operational state: Up
LED state: Green On
Maintenance: Disabled
AINS Soak: Running
  Total Duration: 0 hour(s) 15 minute(s)
  Remaining Duration: 0 hour(s) 5 minute(s) 37 second(s)
Laser Squelch: Disabled

```

```

Phy:
  Media type: Not known

```

Autonegotiation disabled.

```

Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: None (or external)
  BER monitoring:
    Not supported
  Holdoff Time: 0ms

```

This example displays the ethernet controller statistics with AINS Soak in pending state.

```

RP/0/RP0/CPU0:ios#show controllers HuC 0/0/0/2
Thu Mar 12 13:52:12.129 UTC
Operational data for interface HundredGigECtrlr0/0/0/2:

```

```

State:
  Administrative state: enabled
  Operational state: Down (Reason: State undefined)
  LED state: Red On
  Maintenance: Disabled
  AINS Soak: Pending
    Total Duration: 0 hour(s) 30 minute(s)
    Remaining Duration: 0 hour(s) 30 minute(s) 0 second(s)
  Laser Squelch: Disabled

```

```

Phy:
  Media type: Not known
  Alarms:
    Current:
      Local Fault
  Statistics:
    FEC:
      Corrected Codeword Count: 0
      Uncorrected Codeword Count: 9

```

Autonegotiation disabled.

```

Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: None (or external)
  BER monitoring:
    Not supported
  Forward error correction: Standard (Reed-Solomon)
  Holdoff Time: 0ms

```

This example displays the optics controller statistics with AINS Soak in running state.

```

RP/0/RP0/CPU0:ios#show controller optics 0/1/0/3

```

Thu Feb 21 19:45:41.088 UTC

Controller State: Up

Transport Admin State: Automatic In Service

Laser State: On

LED State: Green

Optics Status

Optics Type: Grey optics

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
    
```

Performance Monitoring: Enable

THRESHOLD VALUES

Parameter	High Alarm	Low Alarm	High Warning	Low Warning
Rx Power Threshold(dBm)	4.9	-12.0	0.0	0.0
Tx Power Threshold(dBm)	3.5	-10.1	0.0	0.0
LBC Threshold(mA)	N/A	N/A	0.00	0.00

LBC High Threshold = 98 %
Polarization parameters not supported by optics

Total TX Power = 6.39 dBm

Total RX Power = 5.85 dBm

Lane	Laser Bias	TX Power	RX Power	Output Frequency
1	75.0 %	0.59 dBm	0.63 dBm	230.43 THz
2	68.6 %	0.06 dBm	-0.68 dBm	230.43 THz
3	69.0 %	0.26 dBm	-0.63 dBm	230.43 THz
4	69.1 %	0.56 dBm	-0.10 dBm	230.43 THz

Transceiver Vendor Details

```

Form Factor      : QSFP28
Name             : CISCO-FINISAR
Part Number      : FTLC1152RGPL-C2
Rev Number       : CISCO-FINISAR
Serial Number    : FNS22150LEC
    
```

```

PID                : QSFP-100G-CWDM4-S
VID                : V02
CISCO-FINISAR
Date Code(yy/mm/dd) : 18/04/11
Fiber Connector Type: LC
Sonet Application Code: Not Set
Ethernet Compliance Code: 100GBASE-CWDM4

```

Transceiver Temperature : 32 Celsius

```

AINS Soak           : Running
AINS Timer        : 0h, 15m
AINS remaining time : 771 seconds

```

When the soak time expires, AINS state changes from Running to None. The Transport Admin State of optics controller changes from Automatic In Service to In Service.

```
RP/0/RP0/CPU0:ios# show controllers optics 0/1/0/3
```

Thu Feb 21 20:02:34.126 UTC

Controller State: Up

Transport Admin State: In Service

Laser State: On

LED State: Green

Optics Status

Optics Type: Grey optics

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

```

Performance Monitoring: Enable

THRESHOLD VALUES

Parameter	High Alarm	Low Alarm	High Warning	Low Warning
Rx Power Threshold(dBm)	4.9	-12.0	0.0	0.0
Tx Power Threshold(dBm)	3.5	-10.1	0.0	0.0
LBC Threshold(mA)	N/A	N/A	0.00	0.00

```
LBC High Threshold = 98 %
Polarization parameters not supported by optics
```

```
Total TX Power = 6.41 dBm
```

```
Total RX Power = 5.85 dBm
```

Lane	Laser Bias	TX Power	RX Power	Output Frequency
1	74.9 %	0.60 dBm	0.63 dBm	230.43 THz
2	68.6 %	0.06 dBm	-0.70 dBm	230.43 THz
3	69.0 %	0.30 dBm	-0.63 dBm	230.43 THz
4	69.1 %	0.57 dBm	-0.11 dBm	230.43 THz

Transceiver Vendor Details

```
Form Factor           : QSFP28
Name                  : CISCO-FINISAR
Part Number           : FTLC1152RGPL-C2
Rev Number            : CISCO-FINISAR
Serial Number         : FNS22150LEC
PID                   : QSFP-100G-CWDM4-S
VID                   : V02
CISCO-FINISAR
Date Code(yy/mm/dd)  : 18/04/11
Fiber Connector Type: LC
Sonet Application Code: Not Set
Ethernet Compliance Code: 100GBASE-CWDM4
```

```
Transceiver Temperature : 32 Celsius
```

```
AINS Soak           : None
AINS Timer          : 0h, 0m
AINS remaining time : 0 seconds
```

Configuring AINS on OTN-XP Card

You can configure the default AINS settings for all controllers on the OTN-XP card using the shared plane configuration. The configuration is applied to any line card that is installed in the NCS 1004. Use the following commands:

configure

```
ains-soak hours minutes minutes
```

commit

The following is a sample in which all the controllers on the OTN-XP card are configured with AINS with soak time period specified to be two minutes.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#ains-soak hours 0 minutes 2
RP/0/RP0/CPU0:ios(config)#commit
RP/0/RP0/CPU0:ios(config)#do show controllers optics 0/1/0/0
Tue Apr 28 11:50:15.431 UTC

Controller State: Down

Transport Admin State: Automatic In Service
```

Laser State: On

LED State: Red

Optics Status

Optics Type: 100G QSFP28 LR4

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

Performance Monitoring: Enable

THRESHOLD VALUES

Parameter	High Alarm	Low Alarm	High Warning	Low Warning
Rx Power Threshold(dBm)	4.9	-12.0	0.0	0.0
Tx Power Threshold(dBm)	3.5	-10.1	0.0	0.0
LBC Threshold(mA)	N/A	N/A	0.00	0.00

LBC High Threshold = 98 %

Polarization parameters not supported by optics

Total TX Power = 7.74 dBm

Total RX Power = -40.00 dBm

Lane	Laser Bias	TX Power	RX Power	Output Frequency
1	67.2 %	1.85 dBm	-40.00 dBm	231.39 THz
2	67.9 %	1.55 dBm	-40.00 dBm	230.59 THz
3	67.5 %	1.58 dBm	-40.00 dBm	229.79 THz
4	66.8 %	1.89 dBm	-40.00 dBm	230.25 THz

Transceiver Vendor Details

Form Factor : QSFP28
Name : CISCO-FINISAR
Part Number : 10-3204-01
Rev Number : B
Serial Number : FNS20510YUB
PID : ONS-QSFP28-LR4
VID : V01
Date Code (yy/mm/dd) : 16/12/15
Fiber Connector Type: LC
Otn Application Code: 4I1-9D1F
Sonet Application Code: Not Set

```

Ethernet Compliance Code: 100GBASE-LR4

Transceiver Temperature : 27 Celsius

```

```

AINS Soak           : Pending
AINS Timer          : 0h, 2m
AINS remaining time : 120 seconds

```

To override the default AINS settings on a specific controller, use the following commands:

automatic-in-service controller optics *R/S/I/P hours minutes minutes*



Note This configuration does not persist after an RP reload operation.

The following is a sample in which the optics controller on the OTN-XP card is configured with a soak time period of 45 minutes.

```

RP/0/RP0/CPU0:ios#automatic-in-service controller optics 0/1/0/0 hours 0 minutes 45
Tue Apr 28 11:55:15.666 UTC
RP/0/RP0/CPU0:ios#show controllers optics 0/1/0/0
Tue Apr 28 11:55:30.323 UTC

```

Controller State: Down

Transport Admin State: Automatic In Service

Laser State: On

LED State: Red

Optics Status

Optics Type: 100G QSFP28 LR4

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

```

Performance Monitoring: Enable

THRESHOLD VALUES

Parameter	High Alarm	Low Alarm	High Warning	Low Warning
-----------	------------	-----------	--------------	-------------

```

-----
Rx Power Threshold(dBm)      4.9      -12.0      0.0      0.0
Tx Power Threshold(dBm)     3.5      -10.1      0.0      0.0
LBC Threshold(mA)           N/A      N/A      0.00     0.00

```

```

LBC High Threshold = 98 %
Polarization parameters not supported by optics

```

```
Total TX Power = 7.74 dBm
```

```
Total RX Power = -40.00 dBm
```

```

Lane   Laser Bias   TX Power   RX Power   Output Frequency
-----
1      67.2 %      1.85 dBm  -40.00 dBm  231.39 THz
2      67.9 %      1.55 dBm  -40.00 dBm  230.59 THz
3      67.5 %      1.58 dBm  -40.00 dBm  229.79 THz
4      66.8 %      1.89 dBm  -40.00 dBm  230.25 THz

```

Transceiver Vendor Details

```

Form Factor      : QSFP28
Name             : CISCO-FINISAR
Part Number      : 10-3204-01
Rev Number       : B
Serial Number    : FNS20510YUB
PID              : ONS-QSFP28-LR4
VID              : V01
Date Code(yy/mm/dd) : 16/12/15
Fiber Connector Type: LC
Otn Application Code: 4I1-9D1F
Sonet Application Code: Not Set
Ethernet Compliance Code: 100GBASE-LR4

```

```
Transceiver Temperature : 27 Celsius
```

```

AINS Soak          : Pending
AINS Timer       : 0h, 45m
AINS remaining time : 2700 seconds

```

FEC

Forward error correction (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.

FEC States for Ethernet Controller

The following table lists the FEC states for the Ethernet controller.

State	Description
None	FEC is not enabled on the Ethernet controller.

State	Description
Standard	Standard (Reed-Solomon) FEC is enabled on the Ethernet controller.

FEC configuration is automatically enabled for only the pluggables that support Auto-FEC. If you manually configure FEC, the manual configuration overrides the Auto-FEC.

The supported pluggables for Auto-FEC are:

- QSFP-100G-SR4-S
- QSFP-100G-CWDM4-S
- QSFP-100G-SM-SR
- QSFP-100G-AOC-1M
- QSFP-100G-AOC-3M
- QSFP-100G-AOC-10M
- QDD-400-AOC15M
- QDD-400G-FR4-S
- QSFP-100G-ER4L
- QDD-400G-DR4-S
- QDD-400G-LR8-S

The LR4 pluggable is a 1310nm long range band pluggable that does not require you to enable FEC.

The software automatically enables FEC mode on the pluggables installed in the Cisco NCS 1004. When you upgrade the software of an NCS 1004 with pluggables in the FEC disabled mode, traffic is affected.

The following sample shows the running FEC configuration on the LR4 pluggable:

```
RP/0/RP0/CPU0:ios#show controller HundredGigEctrler 0/0/0/4
Thu Aug  8 15:41:20.857 IST
Operational data for interface HundredGigEctrler0/0/0/4:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Disabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Enabled

Phy:
  Media type: Not known

Autonegotiation disabled.

Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
```

```

Flowcontrol: None
Loopback: None (or external)
BER monitoring:
    Not supported
Holdoff Time: 0ms

```

The following sample shows the running FEC configuration on the non LR4 pluggable:

```

RP/0/RP0/CPU0:ios#show controller HundredGigECtrlr 0/0/0/2
Thu Aug  8 15:41:56.457 IST
Operational data for interface HundredGigECtrlr0/0/0/2:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Disabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Disabled

Phy:
  Media type: Not known
  Statistics:
    FEC:
      Corrected Codeword Count: 0
      Uncorrected Codeword Count: 66

Autonegotiation disabled.

Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: None (or external)
  BER monitoring:
    Not supported
  Forward error correction: Standard (Reed-Solomon)
  Holdoff Time: 0ms

```

Configuring FEC on the Ethernet Controller



Note The FEC configuration is not required for the supported pluggables. The configuration is required only in the case of non-Cisco qualified non-LR4 pluggables.

To configure FEC on the Ethernet controller, use the following command:

```

configure
controller HundredGigECtrlr R/S/I/P fec { none | standard }
commit

```

The following sample shows how to configure FEC on the Ethernet controller:

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller HundredGigEctrlr 0/1/0/10 fec standard
RP/0/RP0/CPU0:ios(config)#commit
```

The following sample shows the running FEC configuration on the Ethernet controller:

```
RP/0/RP0/CPU0:BH-SIT2#show controller HundredGigEctrlr 0/1/0/10
Tue Jul 16 15:30:30.165 IST
Operational data for interface HundredGigEctrlr0/1/0/10:
```

```
State:
  Administrative state: enabled
  Operational state: Down (Reason: State undefined)
  LED state: Red On
  Maintenance: Disabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Disabled
```

```
Phy:
  Media type: Not known
  Alarms:
    Current:
      Loss of Frequency Sync Data
  Statistics:
    FEC:
      Corrected Codeword Count: 0
      Uncorrected Codeword Count: 0
```

Autonegotiation disabled.

```
Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: None (or external)
  BER monitoring:
    Not supported
  Forward error correction: Standard (Reed-Solomon)
  Holdoff Time: 0ms
```

FEC States for CoherentDSP Controller

The following table lists the FEC states for the coherentDSP controllers.

Table 5: FEC State for CoherentDSP Controllers

State	Description
EnhancedSD15	FEC Soft-Decision 15.
EnhancedSD27	FEC Soft-Decision 27. Default.

Configuring FEC on CoherentDSP Controllers

To configure FEC on the CoherentDSP controller, use the following command:

```

configure
controller coherentDSP R/S/I/P
fec {EnhancedSD15 | EnhancedSD27}
commit

```

The following sample shows how to configure FEC on the CoherentDSP controller:

```

RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller coherentDSP 0/0/0/0
RP/0/RP0/CPU0:ios(config-CoDSP)#fec EnhancedSD15
Tue Feb 25 11:25:52.670 UTC
WARNING! Changing FEC mode can impact traffic
RP/0/RP0/CPU0:ios(config-CoDSP)#commit

```

Verifying FEC on CoherentDSP Controllers

The following sample shows the FEC configuration on the CoherentDSP controller:

```

RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/0/0/0

Tue Feb 25 11:26:08.235 UTC

Port                               : CoherentDSP 0/0/0/0
Controller State                    : Up
Inherited Secondary State          : Normal
Configured Secondary State         : Normal
Derived State                       : In Service
Loopback mode                       : None
BER Thresholds                     : SF = 1.0E-5  SD = 1.0E-7
Performance Monitoring              : Enable
Bandwidth                           : 50.0Gb/s
Alarm Information:
LOS = 1 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
Detected Alarms                     : None
Bit Error Rate Information
PREFEC BER                          : 0.00E+00
POSTFEC BER                         : 0.00E+00
Q-Factor                            : 0.00 dB
Q-Margin                             : -5.00dB
Instantaneous Q_margin           : 0 dB

TTI :
Remote IP addr                       : 0.0.0.0
FEC mode                             : Soft-Decision 15

AINS Soak                            : None
AINS Timer                            : 0h, 0m
AINS remaining time                  : 0 seconds

```

Laser Squelching

You can enable laser squelching on Ethernet controllers. When laser squelching is enabled, the laser is shut down in the event of trunk faults (LOS, LOF), and a SQUELCHED alarm is raised on the mapped client port.

To configure laser squelching on the Ethernet controllers, use the following commands:

configure

controller HundredGigECtrlr *Rack/Slot/Instance/Port*

laser-squelch

commit

The following is a sample where laser squelching is enabled on the Ethernet controller.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller HundredGigECtrlr 0/1/0/10
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#laser-squelch
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#commit
```

The following is a sample to view the laser squelch status on the controller.

```
RP/0/RP0/CPU0:ios#show controller HundredGigECtrlr 0/1/0/10
Fri Feb 22 15:18:47.011 UTC
Operational data for interface HundredGigECtrlr0/1/0/10:
```

State:

```
Administrative state: enabled
Operational state: Up
LED state: Green On
Maintenance: Disabled
AINS Soak: None
  Total Duration: 0 hour(s) 0 minute(s)
  Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
Laser Squelch: Enabled
```

Phy:

```
Media type: Not known
Statistics:
  FEC:
    Corrected Codeword Count: 0
    Uncorrected Codeword Count: 0
```

Autonegotiation disabled.

Operational values:

```
Speed: 100Gbps
Duplex: Full Duplex
Flowcontrol: None
Loopback: None (or external)
BER monitoring:
  Not supported
Forward error correction: Standard (Reed-Solomon)
Holdoff Time: 0ms
```

Configuring Laser Squelching on OTN-XP Card

From R7.2.1 onwards, laser squelching is supported on 10GE controllers for the OTN-XP card.

Configuring Laser Squelching on 10GE Controllers

To configure laser squelching on the 10GE controllers for the OTN-XP card, use the following commands:

configure

controller tenGigEctr1r *Rack/Slot/Instance/Port/Lanenumbr*

laser-squelch

commit

The range of *Lanenumbr* is from 1 to 4.

The following is a sample where laser squelching is enabled on the 10GE controller for the OTN-XP card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller tenGigEctr1r 0/0/0/4/1
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#laser-squelch
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#commit
```

The following is a sample to view the laser squelch status on the 10GE controller.

```
P/0/RP0/CPU0:ios#show controllers tenGigEctr1r 0/0/0/4/1
Wed May 6 06:28:29.603 UTC
Operational data for interface TenGigEctr1r0/0/0/4/1:

State:
Administrative state: enabled
Operational state: Up
LED state: Green On
Maintenance: Disabled
AINS Soak: None
Total Duration: 0 hour(s) 0 minute(s)
Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
Laser Squelch: Enabled

Phy:
Media type: Not known

Autonegotiation disabled.

Operational values:
Speed: 10Gbps
Duplex: Full Duplex
Flowcontrol: None
Loopback: None (or external)
Inter-packet gap: standard (12)
BER monitoring:
Not supported
Holdoff Time: 0ms
```

Idle Insertion

When a fault occurs on the trunk port, you can hold the propagation of local faults using the idle insertion feature. This feature is enabled on the ethernet controller by configuring the hold-off timer.

When the fault occurs on the trunk, idles are inserted in the traffic stream from the trunk port to the client port for the duration of the configured holdoff-time. If the trunk port remains faulty beyond the configured holdoff-time, a local fault is transmitted towards the client device. If the trunk recovers from the fault before the holdoff-time expires, traffic resumes.

This feature can be used on customer deployments to prevent reset of client ports during a PSM switchover.

You can enable the idle insertion feature by using the following commands:

configure

controller HundredGigECtrlr *Rack/Slot/Instance/Port*

holdoff-time trunk-fault *time-value*

The range of *timevalue* is from 0 ms to 3000 ms.

The following is a sample for enabling the hold off -timer in 100GE controllers:

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller HundredGigECtrlr 0/1/0/10
RP/0/RP0/CPU0:ios (config-eth-ctrlr)#holdoff-time trunk-fault 3000
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#commit
```

To view the hold-off time that is configured on 100GE controller, use the following command:

show controllers hundredGigECtrlr *Rack/Slot/Instance/Port*

Example

```
RP/0/RP0/CPU0:ios#show controllers HundredGigECtrlr 0/1/0/10
Fri Feb 22 18:58:06.888 UTC
Operational data for interface HundredGigECtrlr0/1/0/10:
```

State:

```
Administrative state: enabled
Operational state: Up
LED state: Green On
Maintenance: Disabled
AINS Soak: None
  Total Duration: 0 hour(s) 0 minute(s)
  Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
Laser Squelch: Disabled
```

Phy:

```
Media type: Not known
Statistics:
  FEC:
    Corrected Codeword Count: 0
    Uncorrected Codeword Count: 0
```

Autonegotiation disabled.

Operational values:

```
Speed: 100Gbps
Duplex: Full Duplex
Flowcontrol: None
Loopback: None (or external)
```

```

BER monitoring:
  Not supported
Forward error correction: Standard (Reed-Solomon)
Holdoff Time: 3000ms

```

Enabling Idle Insertion on OTN-XP Card

From R7.2.1 onwards, you can enable the idle insertion feature on the 10GE controller for OTN-XP card.

To enable idle insertion on the 10GE controller, enter the following commands:

configure

```
controller tenGigECtrlr Rack/Slot/Instance/Port/Lanenum
```

```
holdoff-time trunk-fault time-value
```

commit

The range of *Lanenum* is from 1 to 4 and the range of holdoff-time trunk-fault *time-value* is from 0 to 3000 ms.

The following is a sample for enabling the idle insertion feature in 10GE controllers:

```

RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller tenGigE Ctrlr 0/0/0/4/1
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#holdoff-time trunk-fault 2000
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#commit

```

To view the hold-off time that is configured on 10GE controllers, use the following command:

```
show controllers tenGigECtrlr Rack/Slot/Instance/Port/Lanenum
```

Example

```

RP/0/RP0/CPU0:ios#show controllers TenGigECtrlr 0/0/0/4/1
Thu Mar 26 12:46:16.543 UTC
Operational data for interface TenGigE Ctrlr0/0/0/4/1:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Disabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Disabled

Phy:
  Media type: Not known

Autonegotiation disabled.

Operational values:
  Speed: 10Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: None (or external)
  Inter-packet gap: standard (12)
  BER monitoring:
    Not supported
  Holdoff Time: 2000ms

```

LLDP Drop

Link Layer Discovery Protocol (LLDP) Snooping is enabled by default on all ethernet controllers.

To verify the LLDP neighbors, use the following commands:

```
RP/0/RP0/CPU0:ios#show lldp neighbors detail
Tue Mar 12 11:49:20.819 IST
Capability codes:
    (R) Router, (B) Bridge, (T) Telephone, (C) DOCSIS Cable Device
    (W) WLAN Access Point, (P) Repeater, (S) Station, (O) Other

-----
Local Interface: HundredGigEctrlr0/1/0/7
Chassis id: 008a.96cd.34e1
Port id: Hu0/0/0/4
Port Description - not advertised
System Name: ncs5500_node

System Description:
    6.1.4, NCS-5500

Time remaining: 116 seconds
Hold Time: 120 seconds
System Capabilities: R
Enabled Capabilities: R
Management Addresses - not advertised
Peer MAC Address: 00:8a:96:cd:34:10

-----
Local Interface: HundredGigEctrlr0/1/0/13
Chassis id: 008a.96cd.34e1
Port id: Hu0/0/0/5
Port Description - not advertised
System Name: ncs5500_node

System Description:
    6.1.4, NCS-5500

Time remaining: 90 seconds
Hold Time: 120 seconds
System Capabilities: R
Enabled Capabilities: R
Management Addresses - not advertised
Peer MAC Address: 00:8a:96:cd:34:14

Total entries displayed: 2

RP/0/RP0/CPU0:ios#show lldp neighbors
Tue Mar 12 16:17:56.713 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
ncs5500_node       HundredGigEctrlr0/1/0/7  120       R           Hu0/0/0/4
ncs5500_node       HundredGigEctrlr0/1/0/13 120       R           Hu0/0/0/5

Total entries displayed: 2
```

When you enable LLDP drop on the client controller ports of the muxponder or muxponder slice, the LLDP frames drop on the ports without forwarding.

Configuring LLDP Drop

You can configure the LLDP drop for a muxponder or muxponder slice. By default, the LLDP drop status is set to False. On enabling the LLDP Drop, the status is set to True.

To configure LLDP drop on a muxponder use the following command:

configure

hw-module location *location* mxponder drop-lldp



Note Use the **no** form of the command to disable LLDP drop.

commit

Limitation

- When you disable LLDP globally, the LLDP gets disabled on all the interfaces.



Note By default, LLDP is enabled for NCS 1004. But when you enable and disable LLDP in the global configuration mode, LLDP gets disabled on all the interfaces.

Workaround: You must enable LLDP globally or reload the Router.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios#hw-module location 0/1 mxponder drop-lldp
RP/0/RP0/CPU0:ios#commit
```

configure

hw-module location *location* mxponder-slice *slice-number* drop-lldp



Note Use the **no** form of the command to disable LLDP drop.

To configure LLDP drop on a muxponder slice, use the following command:

commit

The following is a sample in which slice 0 client ports are enabled with LLDP drop.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 0 drop-lldp
RP/0/RP0/CPU0:ios(config)#commit
```

Verifying the Status of LLDP Drop

To verify the LLDP drop enabled status, use the following command.

```
RP/0/RP0/CPU0:ios#show hw-module location all mxponder
Fri Feb 22 13:22:19.281 UTC
```

```
Location:          0/0
Client Bitrate:    NONE
Trunk Bitrate:     NONE
Status:           Not Provisioned
```

```
Location:          0/1
Slice ID:          0
Client Bitrate:    100GE
Trunk Bitrate:     500G
Status:           Provisioned
LLDP Drop Enabled: FALSE
Client Port                Mapper/Trunk Port                CoherentDSP0/1/0/0
                           Traffic Split Percentage
HundredGigECtrlr0/1/0/2    ODU40/1/0/0/0                    100
HundredGigECtrlr0/1/0/3    ODU40/1/0/0/1                    100
HundredGigECtrlr0/1/0/4    ODU40/1/0/0/2                    100
HundredGigECtrlr0/1/0/5    ODU40/1/0/0/3                    100
HundredGigECtrlr0/1/0/6    ODU40/1/0/0/4                    100
```

```
Location:          0/1
Slice ID:          1
Client Bitrate:    100GE
Trunk Bitrate:     500G
Status:           Provisioned
LLDP Drop Enabled: FALSE
Client Port                Mapper/Trunk Port                CoherentDSP0/1/0/1
                           Traffic Split Percentage
HundredGigECtrlr0/1/0/8    ODU40/1/0/1/0                    100
HundredGigECtrlr0/1/0/9    ODU40/1/0/1/1                    100
HundredGigECtrlr0/1/0/10   ODU40/1/0/1/2                    100
HundredGigECtrlr0/1/0/11   ODU40/1/0/1/3                    100
HundredGigECtrlr0/1/0/12   ODU40/1/0/1/4                    100
```

```
Location:          0/2
Slice ID:          0
Client Bitrate:    100GE
Trunk Bitrate:     500G
Status:           Provisioned
LLDP Drop Enabled: FALSE
Client Port                Mapper/Trunk Port                CoherentDSP0/2/0/0
                           Traffic Split Percentage
HundredGigECtrlr0/2/0/2    ODU40/2/0/0/0                    100
HundredGigECtrlr0/2/0/3    ODU40/2/0/0/1                    100
HundredGigECtrlr0/2/0/4    ODU40/2/0/0/2                    100
HundredGigECtrlr0/2/0/5    ODU40/2/0/0/3                    100
HundredGigECtrlr0/2/0/6    ODU40/2/0/0/4                    100
```

```
Location:          0/2
Slice ID:          1
Client Bitrate:    100GE
Trunk Bitrate:     500G
Status:           Provisioned
LLDP Drop Enabled: FALSE
Client Port                Mapper/Trunk Port                CoherentDSP0/2/0/1
```

		Traffic Split Percentage	
HundredGigECtrlr0/2/0/8	ODU40/2/0/1/0		100
HundredGigECtrlr0/2/0/9	ODU40/2/0/1/1		100
HundredGigECtrlr0/2/0/10	ODU40/2/0/1/2		100
HundredGigECtrlr0/2/0/11	ODU40/2/0/1/3		100
HundredGigECtrlr0/2/0/12	ODU40/2/0/1/4		100
Location:	0/3		
Slice ID:	0		
Client Bitrate:	100GE		
Trunk Bitrate:	300G		
Status:	Provisioned		
LLDP Drop Enabled:	TRUE		
Client Port	Mapper/Trunk Port	CoherentDSP0/3/0/0	
	Traffic Split Percentage		
HundredGigECtrlr0/3/0/2	ODU40/3/0/0/0		100
HundredGigECtrlr0/3/0/3	ODU40/3/0/0/1		100
HundredGigECtrlr0/3/0/4	ODU40/3/0/0/2		100

Link Layer Discovery Protocol (LLDP) Support on Management Interface

The LLDP support on management interface feature requires a system to form LLDP neighborship over the system management interface, through which it advertises and learns LLDP neighbor information. This information about neighbors used to learn about the neighbors and in turn the topology of the devices for Operations, Administration, and Maintenance (OAM) purposes.

Advantages of LLDP

- Provides support on non-Cisco devices.
- Enables neighbor discovery between non-Cisco devices.

Limitation

- When you disable LLDP globally, the LLDP gets disabled on all the interfaces.



Note By default, LLDP is enabled for NCS 1004. But when you enable and disable LLDP in the global configuration mode, LLDP gets disabled on all the interfaces.

Workaround: You must enable LLDP globally or reload the Router.

Cisco Discovery Protocol (CDP) vs LLDP

The CDP is a device discovery protocol that runs over Layer 2. Layer 2 is also known as the data link layer that runs on all Cisco devices, such as routers, bridges, access servers, and switches. This protocol allows the network management applications to automatically discover and learn about other Cisco devices that connect to the network.

The LLDP is also a device discovery protocol that runs over Layer 2. This protocol allows the network management applications to automatically discover and learn about other non-Cisco devices that connect to the network.

Interoperability between non-Cisco devices using LLDP

LLDP is also a neighbor discovery protocol that is used by 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.

With LLDP, the user can also access the information about a particular physical network connection. If the user uses a non-Cisco monitoring tool (through SNMP), LLDP helps you identify the Object Identifiers (OIDs) that the system supports. The following OIDs are supported:

- 1.0.8802.1.1.2.1.4.1.1.4
- 1.0.8802.1.1.2.1.4.1.1.5
- 1.0.8802.1.1.2.1.4.1.1.6
- 1.0.8802.1.1.2.1.4.1.1.7
- 1.0.8802.1.1.2.1.4.1.1.8
- 1.0.8802.1.1.2.1.4.1.1.9
- 1.0.8802.1.1.2.1.4.1.1.10
- 1.0.8802.1.1.2.1.4.1.1.11
- 1.0.8802.1.1.2.1.4.1.1.12

Neighbor Discovery

System advertises the LLDP TLV (Type Length Value) details over the management network using which other devices in the management network can learn about this device.

Configuring LLDP

- LLDP full stack functionality is supported on all three management interfaces supported in NCS 1004.
- You can selectively enable or disable LLDP on any of the management interfaces on demand.
- You can selectively enable or disable LLDP transmit or receive functionality at the management interface level.
- Information gathered using LLDP can be stored in the device Management Information Database (MIB) and queried with the Simple Network Management protocol (SNMP).
- LLDP operational data are available in both Command Line Interface and netconf-yang interface.

Enabling LLDP Globally

When you enable LLDP globally, all interfaces that support LLDP are automatically enabled for both transmit and receive operations.



Note You can override this default operation at the interface to disable receive or transmit operations.

The following table describes the global LLDP attributes that the user can configure:

Table 6:

Attribute	Default	Range	Description
Holdtime	120	0–65535	Specifies the hold time (in sec). Hold time refers to the time or duration that an LLDP device maintains the neighbor information before discarding.
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)

The following example shows the commands to configure LLDP globally. The global LLDP configuration enables LLDP on all the three management interfaces.

```
RP/0/RP0/CPU0:regen#configure terminal
RP/0/RP0/CPU0:regen(config)#lldp management enable
RP/0/RP0/CPU0:regen(config)#lldp holdtime 30
RP/0/RP0/CPU0:regen(config)#lldp reinit 2
RP/0/RP0/CPU0:regen(config)#commit
```

Verification

You can verify the LLDP configuration using the **show running-config lldp** command.

The output of **show running-config lldp** command is as follows:

```
RP/0/RP0/CPU0:regen#show running-config lldp
Tue Dec 10 10:36:11.567 UTC
lldp
timer 30
reinit 2
holdtime 120
management enable
!
```

You can verify the LLDP data using the **show lldp interface** and **show lldp neighbors** commands.

The output of **show lldp interface** command is as follows:

```
RP/0/RP0/CPU0:regen#show lldp interface
Thu Nov 7 08:45:22.934 UTC

MgmtEth0/RP0/CPU0/0:
  Tx: enabled
  Rx: enabled
  Tx state: IDLE
  Rx state: WAIT FOR FRAME

MgmtEth0/RP0/CPU0/1:
```

```
Tx: enabled
Rx: enabled
Tx state: IDLE
Rx state: WAIT FOR FRAME
```

The output of **show lldp neighbors** command is as follows:

```
RP/0/RP0/CPU0:M-131#show lldp neighbors
Mon Dec 2 11:01:20.143 CET
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
[DISABLED]     MgmtEth0/RP0/CPU0/0  120       B           gi19
MYS-130        MgmtEth0/RP0/CPU0/1  120       R           MgmtEth0/RP0/CPU0/1
```

where [DISABLED] shows that the LLDP is disabled on the interface MgmtEth0/RP0/CPU0/0.

Enabling LLDP per Management Interface

The following example shows the commands to configure LLDP at the management interface level.

```
RP/0/RP0/CPU0:ios(config)#interface mgmtEth 0/RP0/CPU0/X
RP/0/RP0/CPU0:ios(config-if)#lldp enable
RP/0/RP0/CPU0:ios(config-if)#commit
```

Disabling LLDP Transmit and Receive Operations

The following example shows the commands to disable the LLDP transmit operations at the specified management interface.

```
RP/0/RP0/CPU0:ios(config)#interface mgmtEth 0/RP0/CPU0/X
RP/0/RP0/CPU0:ios(config-if)#lldp transmit disable
RP/0/RP0/CPU0:ios(config-if)#commit
```

The following example shows the commands to disable the LLDP receive operations at the specified management interface.

```
RP/0/RP0/CPU0:ios(config)#interface mgmtEth 0/RP0/CPU0/X
RP/0/RP0/CPU0:ios(config-if)#lldp receive disable
RP/0/RP0/CPU0:ios(config-if)#commit
```

Debugging LLDP Issues

The following commands are used for debugging issues in the LLDP functionality.

- **show lldp traffic**
- **debug lldp all**
- **debug lldp errors**
- **debug lldp events**
- **debug lldp packets**
- **debug lldp tlvs**
- **debug lldp trace**
- **debug lldp verbose**

MAC Address Snooping on Client Ports

MAC address snooping allows you to learn the MAC address of the neighbor, that is connected to the client ports. You can enable ARP snooping on all client ports and learn the MAC address of neighbors through CLI.

This feature overcomes the limitation, where LLDP (Link Layer Discovery protocol) cannot be enabled in some networks.

Limitations

- When you enable or disable MAC address snooping on any slice, few packets are dropped during configuration.
- Open config interface for enabling or disabling MAC address snooping is not supported.
- SNMP MIB is not supported for the MAC address attribute.



Note When you enable MAC address snooping on client ports, it overrides LLDP.

Configuring MAC Address Snooping on Client Ports

You can configure MAC address or ARP snoop on slice in Muxponder slice mode using the following commands.

configure

hw-module location *location mxponder-slice slice-number*

client-rate 100GE

trunk-rate 600G { 100G | 150G | 200G | 250G | 300G | 350G | 400G | 450G | 500G | 550G | 600G }

arp-snoop

commit

Example

The following is a sample in which, MAC address or ARP snoop is configured on the client ports of slice 0 in Muxponder slice mode.

```
RP/0/RP0/CPU0:ios#configure
Mon Mar 16 19:30:33.933 UTC
RP/0/RP0/CPU0:ios(config)#hw-module location 0/3 mxponder-slice 0
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#client-rate 100GE
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#trunk-rate 600G
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#arp-snoop
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#commit
Mon Mar 16 19:30:52.636 UTC
RP/0/RP0/CPU0:ios(config-hwmod-mxp)#end
```

The following is a sample in which, MAC address or ARP snoop is configured in Muxponder mode.

```
RP/0/RP0/CPU0:ios#configure
Mon Mar 16 19:08:17.154 UTC
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder arp-snoop
RP/0/RP0/CPU0:ios(config)#commit
```

The following sample shows the output of **show controllers hundredGigEctr1r** command, before configuring MAC address or ARP snoop on client ports.

```
RP/0/RP0/CPU0:ios#show controllers hundredGigEctr1r 0/1/0/2
Mon Mar 16 19:40:37.434 UTC
Operational data for interface HundredGigEctr1r0/1/0/2:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Disabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Disabled

Phy:
  Media type: Not known

Autonegotiation disabled.

Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: None (or external)
  BER monitoring:
    Not supported
  Holdoff Time: 0ms
```

Viewing Neighbor MAC Address

You can view the neighbor's physical address after enabling MAC address or ARP snoop using the following command. MAC address snoop output is enabled after ARP packets are received on the respective 100G client.

show controllers hundredGigEctr1r R/S/I/P

The following sample shows the neighbor's MAC address after configuring MAC address or ARP snoop on client ports.

```
RP/0/RP0/CPU0:ios#show controllers hundredGigEctr1r 0/1/0/2
Mon Mar 16 19:41:08.047 UTC
Operational data for interface HundredGigEctr1r0/1/0/2:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Disabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Disabled
  Neighbor Address:
    0010.9400.5502
```

```
Phy:
  Media type: Not known

Autonegotiation disabled.

Operational values:
  Speed: 100Gbps
  Duplex: Full Duplex
  Flowcontrol: None
```

Loopback

You can configure the loopback on the CoherentDSP, FC, OTU, and Ethernet controllers to identify connection problems. The loopback can be configured only in the maintenance mode. Use the **controller controller-type** and the **secondary-admin-state maintenance** commands to place the controllers in the maintenance mode.



Note Internal and line loopbacks are supported on the FC, OTU, and Ethernet controllers whereas only internal loopbacks are supported on the CoherentDSP controllers.

Configuring Loopback on the 1.2T Card

To configure the loopback, use the following commands:

```
configure
controller controllertype Rack/Slot/Instance/Port
sec-admin-state maintenance
loopback [ line | internal ]
commit
```

Example 1

The following example shows how a line loopback is configured on the Ethernet controller.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller HundredGigEctrler 1/0/1/10 secondary-admin-state
maintenance
RP/0/RP0/CPU0:ios(config)#commit
Fri Feb 22 19:49:46.504 UTC
RP/0/RP0/CPU0:ios(config)#exit
```

The following example shows how to verify a line loopback configured on the Ethernet controller.

```
RP/0/RP0/CPU0:ios#show controller HundredGigEctrler 0/1/0/10
Fri Feb 22 19:50:08.328 UTC
Operational data for interface HundredGigEctrler0/1/0/10:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Enabled
  AINS Soak: Pending
```

```

Total Duration: 0 hour(s) 30 minute(s)
Remaining Duration: 0 hour(s) 30 minute(s) 0 second(s)
Laser Squelch: Enabled

```

Phy:

```

Media type: Not known
Statistics:
  FEC:
    Corrected Codeword Count: 0
    Uncorrected Codeword Count: 0

```

Autonegotiation disabled.

Operational values:

```

Speed: 100Gbps
Duplex: Full Duplex
Flowcontrol: None
Loopback: None (or external)
BER monitoring:
  Not supported
Forward error correction: Standard (Reed-Solomon)
Holdoff Time: 0ms

```

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

```
RP/0/RP0/CPU0:ios(config)#controller HundredGigECtrlr 0/1/0/10 loopback line
```

```
RP/0/RP0/CPU0:ios(config)#commit
```

```
RP/0/RP0/CPU0:ios(config)#exit
```

```
RP/0/RP0/CPU0:ios#show controller HundredGigECtrlr 0/1/0/10
```

```
Fri Feb 22 20:01:00.521 UTC
```

```
Operational data for interface HundredGigECtrlr0/1/0/10:
```

State:

```

Administrative state: enabled
Operational state: Up
LED state: Green On
Maintenance: Enabled
AINS Soak: Pending
  Total Duration: 0 hour(s) 30 minute(s)
  Remaining Duration: 0 hour(s) 30 minute(s) 0 second(s)
Laser Squelch: Enabled

```

Phy:

```

Media type: Not known
Statistics:
  FEC:
    Corrected Codeword Count: 0
    Uncorrected Codeword Count: 6

```

Autonegotiation disabled.

Operational values:

```

Speed: 100Gbps
Duplex: Full Duplex
Flowcontrol: None
Loopback: Line
BER monitoring:
  Not supported
Forward error correction: Standard (Reed-Solomon)
Holdoff Time: 0ms

```

Example 2

The following example shows how to verify an internal loopback configured on the coherent DSP controller.

```

RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/0/0/0
Fri Mar 13 22:00:20.951 UTC

Port                               : CoherentDSP 0/0/0/0
Controller State                    : Up
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                           : 200.0Gb/s

Alarm Information:
LOS = 0 LOF = 1 LOM = 0
OOF = 0 OOM = 0 AIS = 0
IAE = 0 BIAE = 0          SF_BER = 0
SD_BER = 0          BDI = 3 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0
Detected Alarms                : None

Bit Error Rate Information
PREFEC BER                      : 0.00E+00
POSTFEC BER                     : 0.00E+00
Q-Factor                        : 16.70 dB

Q-Margin                         : 0.99dB

TTI :
    Remote hostname              : ios
    Remote interface             : CoherentDSP 0/0/0/0
    Remote IP addr               : 0.0.0.0

FEC mode                         : Soft-Decision 27

AINS Soak                        : None
AINS Timer                       : 0h, 0m
AINS remaining time              : 0 seconds

```

Configuring Loopback on OTN-XP Card

From R7.2.1 onwards, OTN-XP card supports loopback on the OTU2, OTU2e, OTU4, 10GE, and CoherentDSP controllers.

From R7.3.2 onwards, OTN-XP card supports loopback on the 100GE and 400GE controllers.

The CoherentDSP controller supports both line and internal.

To configure the loopback on the controllers, use the following commands:

configure

controller *controller type Rack/Slot/Instance/Port/Lane number*

sec-admin-state maintenance

loopback [**line** | **internal**]

commit

The range of *Lane number* is 1–4.

Example 1

The following example shows how an internal loopback is configured on the 10GE controller.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller tenGigECtrlr 0/0/0/5/2
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#sec-admin-state maintenance
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#loopback internal
RP/0/RP0/CPU0:ios(config-eth-ctrlr)#commit
```

The following example shows how to verify an internal loopback configured on the 10GE controller.

```
RP/0/RP0/CPU0:ios#show controllers tenGigECtrlr 0/0/0/5/2
Thu Apr 23 10:47:48.020 UTC
Operational data for interface TenGigECtrlr0/0/0/5/2:

State:
  Administrative state: enabled
  Operational state: Up
  LED state: Green On
  Maintenance: Enabled
  AINS Soak: None
    Total Duration: 0 hour(s) 0 minute(s)
    Remaining Duration: 0 hour(s) 0 minute(s) 0 second(s)
  Laser Squelch: Disabled

Phy:
  Media type: Not known

Autonegotiation disabled.

Operational values:
  Speed: 10Gbps
  Duplex: Full Duplex
  Flowcontrol: None
  Loopback: Internal
  Inter-packet gap: standard (12)
  BER monitoring:
    Not supported
  Holdoff Time: 0ms
```

Example 2

The following example shows how a line loopback is configured on the OTU2e controller.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller otu2e 0/0/0/11/3
RP/0/RP0/CPU0:ios(config-otu2e)#secondary-admin-state maintenance
RP/0/RP0/CPU0:ios(config-otu2e)#loopback line
RP/0/RP0/CPU0:ios(config-otu2e)#commit
Thu Apr 23 10:55:19.319 UTC
RP/0/RP0/CPU0:ios(config-otu2e)#end
```

The following example shows how to verify a line loopback configured on the OTU2e controller.

```
RP/0/RP0/CPU0:ios#show controllers otu2e 0/0/0/11/3
Thu Apr 23 10:55:28.014 UTC

Port                : OTU2E 0/0/0/11/3
Controller State    : Up
Inherited Secondary State : Normal
Configured Secondary State : Maintenance
Derived State       : Maintenance
Loopback mode       : Line
BER Thresholds      : SF = 1.0E-5  SD = 1.0E-7
```

```

Performance Monitoring           : Enable
Bandwidth                       : 10.0Gb/s

Alarm Information:
LOS = 0 LOF = 1 LOM = 0
OOF = 1 OOM = 1 AIS = 0
IAE = 0 BIAE = 0           SF_BER = 0
SD_BER = 0           BDI = 0 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0
Detected Alarms                 : None

Bit Error Rate Information
PREFEC BER                      : 0.00E+00
POSTFEC BER                     : 0.00E+00

TTI :
    Remote hostname             : ios
    Remote interface           : OTU2E 0/0/0/11/3
    Remote IP addr              : 0.0.0.0

FEC mode                         : STANDARD

AINS Soak                       : None
AINS Timer                      : 0h, 0m
AINS remaining time             : 0 seconds

```

Example 3

The following example shows how an internal loopback is configured on the OTU2 controller.

```

RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller otu2 0/0/0/5/1
RP/0/RP0/CPU0:ios(config-otu2)#secondary-admin-state maintenance
RP/0/RP0/CPU0:ios(config-otu2)#loopback internal
RP/0/RP0/CPU0:ios(config-otu2)#commit
Thu Apr 23 11:01:00.562 UTC
RP/0/RP0/CPU0:ios(config-otu2)#end

```

The following example shows how to verify an internal loopback configured on the OTU2 controller.

```

RP/0/RP0/CPU0:ios#show controllers otu2 0/0/0/5/1
Thu Apr 23 11:01:04.126 UTC

Port                             : OTU2 0/0/0/5/1
Controller State                 : Up
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                       : 10.0Gb/s

Alarm Information:
LOS = 0 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
Detected Alarms                 : None

Bit Error Rate Information

```

```

PREFEC BER                : 0.00E+00
POSTFEC BER               : 0.00E+00

TTI :
    Remote hostname       : SM-TRC SAPI-SECSM-TRC DA
    Remote IP addr       : 209.165.200.229

FEC mode                  : STANDARD

AINS Soak                 : None
AINS Timer                : 0h, 0m
AINS remaining time      : 0 seconds

```

Example 4

The following example shows how an internal loopback is configured on the OTU4 controller.

```

RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller otu4 0/0/0/0
RP/0/RP0/CPU0:ios(config-otu4)#secondary-admin-state maintenance
RP/0/RP0/CPU0:ios(config-otu4)#loopback internal
RP/0/RP0/CPU0:ios(config-otu4)#commit
Thu Apr 23 11:05:22.429 UTC
RP/0/RP0/CPU0:ios(config-otu4)#end

```

The following example shows how to verify an internal loopback configured on the OTU4 controller.

```

RP/0/RP0/CPU0:ios#show controllers otu4 0/0/0/0
Thu Apr 23 11:05:30.281 UTC

Port                       : OTU4 0/0/0/0
Controller State           : Up
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                  : 100.0Gb/s

Alarm Information:
LOS = 1 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
Detected Alarms           : None

Bit Error Rate Information
PREFEC BER                : 0.00E+00
POSTFEC BER               : 0.00E+00

TTI :
    Remote hostname       : ios
    Remote interface     : OTU4 0/0/0/0
    Remote IP addr       : 0.0.0.0

FEC mode                  : STANDARD

AINS Soak                 : None
AINS Timer                : 0h, 0m
AINS remaining time      : 0 seconds

```

Restore Factory Settings



Note Perform this operation only on the console port.

You can restore the factory settings on the NCS 1004. The entire system configuration, including usernames, passwords, and IP addresses, is removed. You can perform this operation only through the console port and not on the management interface. To restore NCS 1004 to factory settings, use the **commit replace** command. After the **commit replace** operation completes, you must perform the IOS XR reload operation.

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.

Example

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#commit replace
Tue Sep 24 09:36:59.430 UTC
```

This commit will replace or remove the entire running configuration. This operation can be service affecting.

Do you wish to proceed? [no]: yes

```
RP/0/RP0/CPU0:ios(config)#exit
```

```
RP/0/RP0/CPU0:ios#reload
```

```
Tue Sep 24 09:38:12.881 UTC
```

```
Standby card not present or not Ready for failover. Proceed? [confirm]
```

```
Preparing system for backup. This may take a few minutes especially for large configurations.
```

```
Status report: node0_RP0_CPU0: BACKUP INPROGRESS
```

```
Status report: node0_RP0_CPU0: BACKUP HAS COMPLETED SUCCESSFULLY
```

```
[Done]
```

```
Proceed with reload? [confirm]
```

```
Reloading node 0/RP0/CPU0
```

```
RL: Reboot initiated with code 1, cause User initiated graceful reload reboot_timeout 30
shutdown delay 0
```

```
RL: Shutdown initiated
```

```
Query the node to be reloaded
```

```
  NODE_IP of noded to be reloaded 192.0.2.4
```

```
sending stop hb
```

```
Cause: User initiated graceful reload
```

```
VM IP addr sent for reload 192.0.2.4
```

```
Received ack from sdrmgr for reload request.Returncode:0
```

```
successful disconnection from service
```

```
wd_disconnect_cb 548 CMP-WD disconnected successfully
```

```
Invmgr successful disconnection from service
```

```
RP/0/RP0/CPU0:ios#
```

```
Disconnecting from 'default-sdr--1' console. Continue(Y/N)?
```

```
Connecting to 'default-sdr--1' console
```




Note The *tti-string* can have a maximum of 64 characters.

The following sample displays how to configure TTI on a coherent DSP controller with the sent and expected strings set to the same ASCII string. The state of the controller is up.

```
RP/0/RP0/CPU0:ios#config
Fri Mar 15 08:03:02.094 UTC
RP/0/RP0/CPU0:ios(config)#controller coherentDSP 0/1/0/1 tti sent ascii 1234
RP/0/RP0/CPU0:ios(config)#controller coherentDSP 0/1/0/1 tti expected ascii 1234
RP/0/RP0/CPU0:ios(config)#commit
Fri Mar 15 08:03:49.725 UTC
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/1/0/1
Fri Mar 15 08:04:06.290 UTC

Port                               : CoherentDSP 0/1/0/1
Controller State                   : Up
Inherited Secondary State         : Normal
Configured Secondary State       : Normal
Derived State                     : In Service
Loopback mode                     : None
BER Thresholds                   : SF = 1.0E-5  SD = 1.0E-7
Performance Monitoring            : Enable

Alarm Information:
LOS = 0 LOF = 0 LOM = 0
OOF = 0 OOM = 0 AIS = 0
IAE = 0 BIAE = 0          SF_BER = 0
SD_BER = 0      BDI = 1 TIM = 0
FECMISMATCH = 0 FEC-UNC = 0
Detected Alarms                 : None

Bit Error Rate Information
PREFEC BER                      : 7.7E-03
POSTFEC BER                     : 0.0E+00

OTU TTI Sent
  OPERATOR SPECIFIC  ASCII      : 1234
  :
  OPERATOR SPECIFIC  HEX        : 31323334000000000000000000000000
  : 00000000000000000000000000000000

OTU TTI Received
  OPERATOR SPECIFIC  ASCII      : 1234
  :
  OPERATOR SPECIFIC  HEX        : 31323334000000000000000000000000
  : 00000000000000000000000000000000

OTU TTI Expected
  OPERATOR SPECIFIC  ASCII      : 1234
  :
  OPERATOR SPECIFIC  HEX        : 31323334000000000000000000000000
  : 00000000000000000000000000000000

FEC mode                          : Soft-Decision 27

AINS Soak                        : None
AINS Timer                       : 0h, 0m
AINS remaining time              : 0 seconds
```

The following example shows how to configure TTI on a coherent DSP controller with the sent and expected strings set to different ASCII strings. The state of the controller goes down and the TIM alarm is raised.

```

RP/0/RP0/CPU0:ios#config
Fri Mar 15 08:54:29.780 UTC
RP/0/RP0/CPU0:ios(config)#controller coherentDSP 0/1/0/1 tti sent ascii 1234
RP/0/RP0/CPU0:ios(config)#controller coherentDSP 0/1/0/1 tti expected ascii 5678
RP/0/RP0/CPU0:ios(config)#commit
Fri Mar 15 08:56:12.293 UTC
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/1/0/1
Fri Mar 15 08:56:33.910 UTC

Port                               : CoherentDSP 0/1/0/1
Controller State                    : Down
Inherited Secondary State          : Normal
Configured Secondary State         : Normal
Derived State                       : In Service
Loopback mode                       : None
BER Thresholds                      : SF = 1.0E-5  SD = 1.0E-7
Performance Monitoring              : Enable

Alarm Information:
LOS = 1 LOF = 0 LOM = 0
OOF = 0 OOM = 0 AIS = 0
IAE = 0 BIAE = 0          SF_BER = 0
SD_BER = 0          BDI = 3 TIM = 1
FECMISMATCH = 0 FEC-UNC = 0
Detected Alarms                  : BDI TIM

Bit Error Rate Information
PREFEC BER                       : 8.2E-03
POSTFEC BER                       : 0.0E+00

OTU TTI Sent
  OPERATOR SPECIFIC ASCII        : 1234
  OPERATOR SPECIFIC HEX          : 31323334000000000000000000000000
  OPERATOR SPECIFIC HEX          : 00000000000000000000000000000000

OTU TTI Received
  OPERATOR SPECIFIC ASCII        : 1234
  OPERATOR SPECIFIC HEX          : 31323334000000000000000000000000
  OPERATOR SPECIFIC HEX          : 00000000000000000000000000000000

OTU TTI Expected
  OPERATOR SPECIFIC ASCII        : 5678
  OPERATOR SPECIFIC HEX          : 35363738000000000000000000000000
  OPERATOR SPECIFIC HEX          : 00000000000000000000000000000000

FEC mode                           : Soft-Decision 27

AINS Soak                          : None
AINS Timer                          : 0h, 0m
AINS remaining time                 : 0 seconds

```

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 the minimum difference between the configured values is equal to or greater than 1500 ps/nm.

The following table lists the default CD search range.

Data Rate	BPS	Card Support	Default CD Search Range
200G to 500G	BPS <= 3	1.2T, 1.2TL	-10,000 to 100,000 ps/nm
	3 < BPS <= 4	1.2T, 1.2TL	-10,000 to 80,000 ps/nm
	4 < BPS <=5	1.2T	-5,000 to 20,000 ps/nm
600G	BPS=5.2578125	1.2T	-2000 to 2,000 ps/nm



Note The cd-min and cd-max values must be set for BPS values that are greater than 4 in the 1.2T card.



Note When the user provisions the cd-min and cd-max values that are outside the range through CLI, the provisioned values are accepted; however, only the actual values supported by the hardware are applied.

The following is a sample where chromatic dispersion is configured on the optics controller.

```
RP/0/RP0/CPU0:ios#configure
Mon Aug 19 19:31:42.115 UTC
RP/0/RP0/CPU0:ios(config)#controller optics 0/1/0/1
RP/0/RP0/CPU0:ios(config-Optics)#cd-max 4000
RP/0/RP0/CPU0:ios(config-Optics)#cd-min -1000
RP/0/RP0/CPU0:ios(config-Optics)#commit
Mon Aug 19 19:35:24.697 UTC
RP/0/RP0/CPU0:ios(config-Optics)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show run controller optics 0/1/0/*
Mon Aug 19 19:57:41.859 UTC
controller Optics0/1/0/0
  transmit-power -15
  dwdm-carrier 50GHz-grid itu-ch 55
  enh-sop-tol-mode 1
  cross-pol-gain-mode 10
  lbc-high-threshold 5
!
controller Optics0/1/0/1
  description trunk power UP
  cd-min -1000
  cd-max 4000
  enh-colorless-mode 2
  enh-sop-tol-mode 3
  nleq-comp-mode 4
  cross-pol-gain-mode 2
  cross-pol-weight-mode 3
  cpr-win-mode 3
  cpr-ext-win-mode 8
  rx-voa fixed-ratio 1200
  filter-roll-off-factor 0.035
!
controller Optics0/1/0/5
  soak-time 10
!
```

Frequency

You can configure the frequency on trunk ports of the line card.

The following table lists the frequency range with grid spacing supported on the line card:

Line Card	Frequency Range (THz)	Default Frequency (THz)	Grid Spacing
1.2T	191.25 to 196.1	193.1	50GHz and 100MHz
1.2TL 1	186.1 to 190.85	188.5	100MHz

¹ Only non-ITU channels are supported

To configure the wavelength, use the following commands:

configure

controller optics *Rack/Slot/Instance/Port*

dwdm-carrier {**100MHz-grid** frequency *frequency*} | {**50GHz-grid** [*frequency frequency*]}

commit

Pseudo Random Binary Sequence

The Pseudo Random Binary Sequence (PRBS) feature allows you to perform data integrity checks between the NCS1004 trunk links without enabling the actual client traffic.

You need to enable PRBS feature on both the transmitting and receiving NCS 1004 trunk ports. The transmitting trunk port generates a bit pattern and sends it to the peer NCS 1004 device. The device detects if the sent bit pattern is received.

You can configure NCS 1004 trunk port in any one of the following modes for PRBS on the 1.2T card:

- **Source mode** — The NCS 1004 at trunk port generates PRBS signal on the line continuously as per the configured PRBS pattern.
- **Sink mode** — The NCS 1004 at trunk port gets locked to the ingress signal according to the configured pattern, analyzes and reports the errors.
- **Source-Sink mode** — The NCS 1004 at trunk port acts as both the PRBS transmitter and receiver, that is, it generates PRBS signal as per the configured pattern, and also gets locked to the ingress signal with the same pattern, and reports the errors.

NCS 1004 trunk port supports the following PRBS patterns:

- **PRBS31** — Sequence length is from $2^{31} - 1$ bits.
- **PRBS23** — Sequence length is from $2^{23} - 1$ bits.
- **PRBS15** — Sequence length is from $2^{15} - 1$ bits.
- **PRBS7** — Sequence length is from $2^7 - 1$ bits.

Limitations of PRBS

There are following limitations with the PRBS feature:

- There is no SNMP support to fetch the PRBS status or Performance Monitoring (PM).
- TTI functionality is not supported with PRBS.
- Loopback and PRBS configurations cannot coexist on a coherentDSP controller. Loopback configuration will be rejected if PRBS is already configured.

PRBS on OTN-XP Card

From R7.2.1 onwards, the OTN-XP card supports PRBS on the mapper optical data unit (ODU2e).



Note ODU2e PRBS is not supported for OTU2E client rates.

NCS 1004 with the OTN-XP card, supports the following PRBS mode:

- **Source mode** — The NCS 1004 at trunk port generates PRBS signal on the line continuously as per the configured PRBS pattern.
- **Sink mode** — The NCS 1004 at trunk port gets locked to the ingress signal according to the configured pattern, analyzes and reports the errors.
- **Source-Sink mode** — The NCS 1004 at trunk port acts as both the PRBS transmitter and receiver, that is, it generates PRBS signal as per the configured pattern, and also gets locked to the ingress signal with the same pattern, and reports the errors.
- **invertedpn31** — Inverted pattern. Sequence length is from $2^{31} - 1$ bits.

NCS 1004 trunk port supports the following PRBS patterns:

- **PRBS31** — Sequence length is from $2^{31} - 1$ bits.
- **PRBS23** — Sequence length is from $2^{23} - 1$ bits.
- **PRBS15** — Sequence length is from $2^{15} - 1$ bits.
- **PRBS7** — Sequence length is from $2^7 - 1$ bits.

Configuring Pseudo Random Binary Sequence

To enable the PRBS on the trunk port, use the following configuration command at the coherentDSP controller:

```
controller coherentDSP R/S//P prbs mode {source | sink | source-sink} pattern {pn31 | pn23 | pn15 | pn7}
```

When the PRBS is enabled on the trunk ports, you can view the following impacts in the corresponding client ports:

- Client traffic is dropped in the direction of source to sink as the frames are overwritten by the PRBS pattern.
- Remote fault is raised on the client ports nearer to the PRBS sink.

Verifying PRBS

R/S/I/P prbs-details

```
RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/1/0/0 prbs-details
Wed Nov 6 23:12:22.464 UTC
```

```
-----PRBS details-----
PRBS Test           : Enable
PRBS Mode           : Source
PRBS Pattern        : PN7
PRBS Status         : Not Applicable
-----
```

- You cannot view any details, if the PRBS is not enabled on the trunk.
- PRBS status is shown as **Not Applicable**, when the mode is **Source**.
- PRBS status is shown as **unlocked**, when the signal is not locked on the receiving side in the **Sink** or **Source-Sink** mode.

Viewing PRBS Performance Monitoring Parameters

PRBS PM parameters are not available for the controllers in Source mode. PRBS PM parameters are reset when PRBS configuration changes on the controller.

To view the PRBS PM parameters on the coherentDSP controller, use the following command:

```
show controllers coherentDSP | ODU4 R/S/I/P pm {current | history } {15-min|24-hour} prbs
```

The following tables describes the fields of PRBS PM parameters.

Table 7: PRBS PM Parameters

PM Parameter	Description
EBC	Cumulative count of PRBS bit errors in the sampling interval (15-minute or 24-hour). PRBS bit errors are accumulated only if PRBS signal is locked.
FOUND-COUNT	Number of state transitions from signal unlocked state to signal locked state in the sampling interval. If state change is not observed in the interval, the count is 0.
LOST-COUNT	Number of state transitions from signal locked state to signal unlocked state in the sampling interval. If state change is not observed in the interval, the count is 0.
FOUND-AT-TS	Latest timestamp when the PRBS state moves from unlocked state to locked state in the sampling interval. If state change is not observed in the interval, the value is null.
CONFIG-PTRN	Configured PRBS pattern on the port.

```
RP/0/RP0:ios#show controllers coherentDSP 0/0/0/1 pm current 15-min prbs
Mon Feb 13 00:58:48.327 UTC

PRBS in the current interval [00:45:00 - 00:58:48 Mon Feb 13 2019]
PRBS current bucket type : Valid
EBC                        : 40437528165
FOUND-COUNT                : 1 FOUND-AT-TS : 00:51:22 Mon Feb 13 2019
LOST-COUNT                 : 1 LOST-AT-TS  : 00:52:52 Mon Feb 13 2019
CONFIG-PTRN                : PRBS_PATTERN_PN31
Last clearing of "show controllers OTU" counters never
```

Configuring PRBS on OTN-XP Card

To configure PRBS mode on the ODU2e controller, you must configure Optical Channel Payload Unit (OPU) on the ODU2e controller followed by the PRBS mode and the pattern. The PRBS supported pattern on the OTN-XP card is invertedPN31.

From R7.3.1 onwards, you can configure PRBS on client or mapper ODU4 and ODU flex controllers.



Note ODU2e PRBS is not supported for OTU2E client rates.

To configure PRBS mode on the ODU2e controller, enter the following commands:

configure

controller *R/S/I/P/client-port/lane-number*

secondary-admin-state maintenance

opu

prbs mode {source | sink | source-sink} **pattern** invertedpn31 {direction {system | line}}

end

commit

The following example shows how to configure PRBS mode as source-sink with pattern as invertedpn31:

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller odu2e0/2/0/12/3/2
RP/0/RP0/CPU0:ios(config-odu2e)#secondary-admin-state maintenance
RP/0/RP0/CPU0:ios(config-odu2e)#opu
RP/0/RP0/CPU0:ios(config-Opuk)#prbs mode source-sink pattern invertedpn31
RP/0/RP0/CPU0:ios(config-Opuk)#end
RP/0/RP0/CPU0:ios(config-odu2e)#commit
```

The following is a sample output of **show controller odu2e** command.

```
RP/0/RP0/CPU0(config-odu2e)#show controller odu2e 0/2/0/12/3/2 prbs-details
Mon Mar 14 21:33:02.293 UTC

-----PRBS details-----
PRBS Test                : Enable
PRBS Mode                : Source-Sink
PRBS Pattern             : INVERTED PN31
PRBS Status              : Locked
PRBS Lock Time(in seconds) : 1190
PRBS Bit Errors          : 0
```

The following example shows how to configure PRBS mode as source-sink with pattern as invertedpn31 with direction as system:

```
RP/0/RP0/CPU0:ios#configure
Wed Nov 11 00:38:11.789 UTC
RP/0/RP0/CPU0:ios(config)#controller odu4 0/2/0/5
RP/0/RP0/CPU0:ios(config-odu4)#secondary-admin-state maintenance
RP/0/RP0/CPU0:ios(config-odu4)#opu prbs mode source-sink pattern invertedpn31 direction
system
RP/0/RP0/CPU0:ios(config-odu4)#commit
Wed Nov 11 00:38:26.391 UTC
```

The following example shows how to configure PRBS mode as source-sink with pattern as invertedpn31 with direction as line:

```
RP/0/RP0/CPU0:ios#configure
Wed Nov 11 00:38:11.789 UTC
RP/0/RP0/CPU0:ios(config)#controller odu4 0/2/0/5
RP/0/RP0/CPU0:ios(config-odu4)#secondary-admin-state maintenance
RP/0/RP0/CPU0:ios(config-odu4)#opu prbs mode source-sink pattern invertedpn31 direction
line
RP/0/RP0/CPU0:ios(config-odu4)#commit
Wed Nov 11 00:38:26.391 UTC
```

Verifying PRBS on OTN-XP Card

You can monitor the status of PRBS on the ODU2e controller using the following command:

show controllers odu2e *R/S/I/P/client-port/client-lane* prbs-details

The following example displays the output of the PRBS configuration with PRBS mode as sink:

```
RP/0/RP0/CPU0:ios#show controllers odu2e 0/2/0/12/3/2 prbs-details
-----PRBS details-----
PRBS Test           : Enable
PRBS Mode           : Sink
PRBS Pattern        : INVERTED PN31
PRBS Status         : Locked
```

The following example displays the output of the PRBS configuration with PRBS mode as source-sink:

```
RP/0/RP0/CPU0:ios#show controllers odu2e 0/2/0/12/3/2 prbs-details
-----PRBS details-----
PRBS Test           : Enable
PRBS Mode           : Source-Sink
PRBS Pattern        : INVERTED PN31
PRBS Status         : Locked
```




CHAPTER 5

Performance Monitoring

- [Performance monitoring parameters, on page 79](#)
- [Clear PM parameters, on page 89](#)

Performance monitoring parameters

A performance monitoring parameter is a network data metric that

- allows service providers to gather, store, set thresholds for, and report network performance statistics,
- supports configurable data collection intervals, such as 30-second, 15-minute, or 24-hour periods, and
- simplifies troubleshooting by enhancing the quality of data collected directly from network equipment.

Service providers can configure and retrieve performance monitoring (PM) counters for various controllers, using intervals that suit operational needs. These parameters provide early detection of network issues, making it easier to maintain the network health and resolve problems quickly.

Instantaneous Q-margins

A Q-margin is a performance monitoring metric that

- measures signal quality in optical transmission systems,
- indicates the margin between observed signal quality and the threshold below which errors increase, and
- allows monitoring and troubleshooting of system stability for coherentDSP controllers on supported cards.

From Release 7.3.1 onwards, instantaneous Q-margin is supported for performance monitoring (PM) parameters on coherentDSP controllers for 1.2T and 1.2TL cards. For more information, see the supporting concept documentation.

In certain scenarios, some initial PM buckets can be displayed as valid, although the instantaneous Q-margin values in those buckets are shown as invalid. PM measurements are typically taken over 30 seconds, 15 minutes, and 24 hours intervals.

The following scenarios can cause initial PM buckets to show valid readings, despite instantaneous Q-margin values being invalid:

- Shutdown or no shutdown on optics
- Bit per second (BPS) change on optics
- Trunk rate change
- Fiber cut

To handle such cases, avoid relying on the initial PM buckets when monitoring instantaneous Q-margin values in these scenarios.

Sample output:

```
RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/2/0/0 pm history flex-bin fec 1
Mon Sep 14 06:16:03.249 UTC

g709 FEC in interval 1 [06:15:50 - 06:16:00 Mon Sep 14 2020]

Flexible bin interval size: 10 seconds
FEC history bucket type : Invalid. ----- > Instantaneous Q_margin is invalid in this
bucket
  EC-BITS   : 38054                               UC-WORDS   : 0

                                MIN                               AVG                               MAX
PreFEC BER      : 0E-15                               3.26E-08                               1.43E-07
PostFEC BER     : 0E-15                               0E-15                                   0E-15
Q               : 0.00                               5.73                                    14.40
Q_margin        : -5.00                               -0.69                                   9.40
Instantaneous Q_margin : -21474836.48                 -8589934.59                             0.00
```

At later intervals, PM buckets may be valid but instantaneous Q-margin values remain invalid, demonstrating why those initial readings should be disregarded.

```
RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/2/0/0 pm history 30-sec fec 1
Mon Sep 14 06:16:53.490 UTC

g709 FEC in interval 1 [06:16:00 - 06:16:30 Mon Sep 14 2020]

FEC history bucket type : Valid ----- > (Instantaneous Q_margin is invalid but the PM
bucket is valid. So these initial bins can be ignored)
  EC-BITS   : 431887                               UC-WORDS   : 0

                                MIN                               AVG                               MAX
PreFEC BER      : 3.97E-09                            4.83E-08                               1.51E-07
PostFEC BER     : 0E-15                               0E-15                                   0E-15
Q               : 14.40                               14.48                                   14.60
Q_margin        : 9.30                               9.46                                    9.60
Instantaneous Q_margin : -21474836.48                 -5010784.19                             14.42
```

Configure PM parameters

This task describes the steps to configure performance monitoring parameters, which are critical for monitoring device operation and analyzing the health of Optics, Ethernet, and coherent DSP controllers.

You can configure and view the performance monitoring parameters for the Optics, Ethernet, and coherent DSP controllers.

Follow these steps to configure PM parameters on Optics, Ethernet, or coherent DSP controllers using CLI commands.

Procedure

- Step 1** Configure the controller using the command **configure controller** *controllertype R/S/I/P*
- Step 2** Configure the pm parameters using the command **pm** { **15-min** | **30-sec** | **24-hour** } { **optics** | **ether** | **pcs** | **fec** | **otn** } { **report** | **threshold** }

Example:

This is a sample in which the performance monitoring parameters of the Optics controller are configured at 24-hour intervals.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller optics 0/0/1/5 pm 24-hour optics threshold osnr max
345
RP/0/RP0/CPU0:ios(config)#commit
```

Example:

This is a sample in which the performance monitoring parameters of the Ethernet controller are configured at 15-minute intervals.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller HundredGigECtrlr 0/3/0/0 pm 15-min pcs report bip
enable
RP/0/RP0/CPU0:ios(config)#commit
```

Example:

This is a sample in which performance monitoring parameters of a Coherent DSP controller are configured at 30-second intervals.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller coherentDSP 0/0/1/1 pm 30-sec fec threshold post-fec-ber
max OE-15
RP/0/RP0/CPU0:ios(config)#commit
```

Upon successful configuration, the specified controllers will begin collecting and retaining PM data according to the set parameters. You will be able to view historical and current performance metrics for troubleshooting and ongoing system health monitoring.

What to do next

Verify PM configuration and review collected PM data using show commands.

View PM parameters

Use this task to view performance monitoring (PM) parameters for Optics, Ethernet, and Coherent DSP controllers. This allows you to assess the operational status, diagnostics, and statistics of controller components for maintenance and troubleshooting.

Performance monitoring parameters provide insight into the health and performance of controller interfaces, including metrics such as optical power, frequency offset, and signal quality. Viewing these parameters helps network operators verify ongoing performance, detect anomalies, and ensure compliance with operational thresholds. Perform this task regularly to support proactive maintenance and rapid fault resolution.

Procedure

Use the command **show controllers *controllertype* *R/S/I/P* { **pm** { **current** | **history** } { **30 sec** | **15-min** | **24-hour** } { **optics** | **ether** | **fec** | **otn** | **prbs** } *linenumber* }**, to view the performance monitoring parameters for Optics, Ethernet, and Coherent DSP controllers.

Example:

This sample displays the current performance monitoring parameters of the Optics controller at 15-minute intervals. Client optics have four lanes.

```
RP/0/RP0/CPU0:ios#show controller optics 0/1/0/3 pm current 15-min optics 3
Sat Feb 9 19:33:42.480 UTC

Optics in the current interval [19:30:00 - 19:33:42 Sat Feb 9 2019]

Optics current bucket type : Valid
      MIN      AVG      MAX      Operational      Configured      TCA      Operational
      Configured      TCA
      Threshold(max) (max)
      Threshold(min)  Threshold(min) (min) Threshold(max)
LBC[% ]      : 0.0      0.0      0.0      0.0      NA      NO      100.0
      NA      NO
OPT[dBm]     : -40.00   -40.00   -40.00   -30.00   NA      NO      63.32
      NA      NO
OPR[dBm]     : -40.00   -40.00   -40.00   -30.00   NA      NO      63.32
      NA      NO
FREQ_OFF[Mhz]: 0      0      0      0      NA      NO      0
      NA      NO
```

This sample displays the current performance monitoring parameters of the Optics controller 15-minute intervals. Trunk optics have one lane.

```
RP/0/RP0/CPU0:ios#show controller optics 0/2/0/1 pm current 15-min optics 1
Sat Feb 9 11:19:15.234 UTC

Optics in the current interval [11:15:00 - 11:19:15 Sat Feb 9 2019]

Optics current bucket type : Valid
      MIN      AVG      MAX      Operational      Configured      TCA      Operational
      Configured      TCA
      Threshold(max) (max)
      Threshold(min)  Threshold(min) (min) Threshold(max)
LBC[% ]      : 0.0      0.0      0.0      0.0      NA      NO      100.0
      NA      NO
OPT[dBm]     : -1.51   -1.49   -1.48   -30.00   NA      NO      63.32
      NA      NO
OPR[dBm]     : -9.11   -9.07   -9.03   -30.00   NA      NO      63.32
      NA      NO
CD[ps/nm]    : 13      15      18      -180000   NA      NO      180000
```

	NA	NO						
DGD[ps]	: 2.00	2.33	3.00	0.01	NA	NO	21474836.46	
	NA	NO						
SOPMD[ps^2]	: 5.00	33.02	79.00	0.01	NA	NO	21474836.46	
	NA	NO						
OSNR[dB]	: 31.50	31.97	32.50	0.01	NA	NO	21474836.46	
	NA	NO						
PDL[dB]	: 0.20	0.34	0.50	0.01	NA	NO	21474836.46	
	NA	NO						
PCR[rad/s]	: 0.00	19.92	93.00	0.01	NA	NO	21474836.46	
	NA	NO						
RX_SIG[dBm]	: -9.05	-9.02	-8.99	-30.00	NA	NO	63.32	
	NA	NO						
FREQ_OFF[Mhz]	: -302	-178	-74	-1500	NA	NO	1500	
	NA	NO						

Performance monitoring parameter outputs for Ethernet, FEC, optics, and Fibre Channel controllers

This reference presents sample outputs from performance monitoring commands for multiple controller types and time intervals. It helps you understand what data is collected and how counters are displayed in operational environments.

This sample displays the current performance monitoring parameters of the Ethernet controller 15-minute intervals.

```
RP/0/RP0/CPU0:ios#show controller HundredGigECtrlr 0/1/0/2 pm current 15-min ether
Fri Aug 30 00:37:53.527 UTC
```

```
ETHER in the current interval [00:30:00 - 00:37:53 Fri Aug 30 2019]
```

```
ETHER current bucket type : Valid
RX-UTIL[%]                : 100.00          Threshold : 0.00          TCA(enable) : NO
TX-UTIL[%]                : 10.00           Threshold : 0.00          TCA(enable) : NO
RX-PKT                    : 3852414442     Threshold : 0           TCA(enable) : NO
STAT-PKT                  : 0              Threshold : 0           TCA(enable) : NO
OCTET-STAT                : 5847965122956 Threshold : 0           TCA(enable) : NO
OVERSIZE-PKT              : 0              Threshold : 0           TCA(enable) : NO
FCS-ERR                   : 0              Threshold : 0           TCA(enable) : NO
LONG-FRAME                : 0              Threshold : 0           TCA(enable) : NO
JABBER-STATS              : 0              Threshold : 0           TCA(enable) : NO
64-OCTET                  : 0              Threshold : 0           TCA(enable) : NO
65-127-OCTET              : 0              Threshold : 0           TCA(enable) : NO
128-255-OCTET             : 0              Threshold : 0           TCA(enable) : NO
256-511-OCTET             : 0              Threshold : 0           TCA(enable) : NO
512-1023-OCTET            : 0              Threshold : 0           TCA(enable) : NO
1024-1518-OCTET           : 0              Threshold : 0           TCA(enable) : NO
IN-UCAST                  : 0              Threshold : 0           TCA(enable) : NO
IN-MCAST                  : 0              Threshold : 0           TCA(enable) : NO
IN-BCAST                  : 0              Threshold : 0           TCA(enable) : NO
OUT-UCAST                 : 0              Threshold : 0           TCA(enable) : NO
OUT-BCAST                 : 0              Threshold : 0           TCA(enable) : NO
```

```

OUT-MCAST           : 0                      Threshold : 0          TCA(enable) : NO
TX-PKT              : 7053588067             Threshold : 0          TCA(enable) : NO
OUT-OCTET           : 451429636288           Threshold : 0          TCA(enable) : NO
IFIN-ERRORS         : 0                      Threshold : 0          TCA(enable) : NO
IFIN-OCTETS         : 0                      Threshold : 0          TCA(enable) : NO
STAT-MULTICAST-PKT  : 0                      Threshold : 0          TCA(enable) : NO
STAT-BROADCAST-PKT  : 0                      Threshold : 0          TCA(enable) : NO
STAT-UNDERSIZED-PKT : 0                      Threshold : 0          TCA(enable) : NO
IN_GOOD_BYTES       : 5847965122956         Threshold : 0          TCA(enable) : NO
IN_GOOD_PKTS        : 3852414442            Threshold : 0          TCA(enable) : NO
IN_DROP_OTHER       : 0                      Threshold : 0          TCA(enable) : NO
OUT_GOOD_BYTES      : 451429636288           Threshold : 0          TCA(enable) : NO
OUT_GOOD_PKTS       : 7053588067            Threshold : 0          TCA(enable) : NO
IN_PKT_64_OCTET     : 0                      Threshold : 0          TCA(enable) : NO
IN_PKTS_65_127_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
IN_PKTS_128_255_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
IN_PKTS_256_511_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
IN_PKTS_512_1023_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
IN_PKTS_1024_1518_OCTETS : 3852414442            Threshold : 0          TCA(enable) : NO
OUT_PKT_64_OCTET    : 7053588067            Threshold : 0          TCA(enable) : NO
OUT_PKTS_65_127_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
OUT_PKTS_128_255_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
OUT_PKTS_256_511_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
OUT_PKTS_512_1023_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
OUT_PKTS_1024_1518_OCTETS : 0                      Threshold : 0          TCA(enable) : NO
TX_UNDERSIZED_PKT   : 0                      Threshold : 0          TCA(enable) : NO
TX_OVERSIZED_PKT    : 0                      Threshold : 0          TCA(enable) : NO
TX_JABBER           : 0                      Threshold : 0          TCA(enable) : NO
TX_BAD_FCS          : 0                      Threshold : 0          TCA(enable) : NO

```



Note Performance monitoring statistics are not supported for IN-UCAST and OUT-UCAST counters for Ethernet clients.

This sample displays the current FEC performance monitoring parameters of the Coherent DSP controller at 15-minute intervals.

```
RP/0/RP0/CPU0:ios#show controller coherentDSP 0/2/0/1 pm current 15-min fec
```

```
Sat Feb 9 11:23:42.196 UTC
```

```
g709 FEC in the current interval [11:15:00 - 11:23:42 Sat Feb 9 2019]
```

```
FEC current bucket type : Valid
```

```

EC-BITS      : 291612035786          Threshold : 903330          TCA(enable) :
YES
UC-WORDS     : 0                      Threshold : 5              TCA(enable) :
YES

```

	MIN	AVG	MAX	Threshold (min)	TCA (enable)	Threshold (max)	TCA (enable)
PreFEC BER :	7.1E-03	7.2E-03	8.1E-03	0E-15	NO	0E-15	NO
PostFEC BER :	0E-15	0E-15	0E-15	0E-15	NO	0E-15	NO

This sample displays the current PRBS performance monitoring parameters of the Coherent DSP controller 15-minute intervals.

```
RP/0/RP0/CPU0:ios#show controllers coherentDSP 0/0/0/1 pm current 15-min prbs
```

```
Mon Feb 13 00:58:48.327 UTC
```

```
PRBS in the current interval [00:45:00 - 00:58:48 Mon Feb 13 2019]
```

```

PRBS current bucket type : Valid
EBC : 40437528165
FOUND-COUNT : 1 FOUND-AT-TS : 00:51:22 Mon Feb 13 2019
LOST-COUNT : 1 LOST-AT-TS : 00:52:52 Mon Feb 13 2019
CONFIG-PTRN : PRBS_PATTERN_PN31
Last clearing of "show controllers OTU" counters never

```

This sample displays the current PCS performance monitoring parameters of the Coherent DSP controller 30-second intervals.

```

RP/0/RP0/CPU0:ios#show controllers hundredGigEctr1r 0/0/0/2 pm current 30-sec pcs
Tue Nov 19 09:17:26.684 UTC

```

```

Ethernet PCS in the current interval [09:17:00 - 09:17:26 Tue Nov 19 2019]

```

```

Ethernet PCS current bucket type : Valid
BIP[00] : 0 Threshold : 0 TCA(enable) : NO
BIP[01] : 0 Threshold : 0 TCA(enable) : NO
BIP[02] : 0 Threshold : 0 TCA(enable) : NO
BIP[03] : 0 Threshold : 0 TCA(enable) : NO
BIP[04] : 0 Threshold : 0 TCA(enable) : NO
BIP[05] : 0 Threshold : 0 TCA(enable) : NO
BIP[06] : 0 Threshold : 0 TCA(enable) : NO
BIP[07] : 0 Threshold : 0 TCA(enable) : NO
BIP[08] : 0 Threshold : 0 TCA(enable) : NO
BIP[09] : 0 Threshold : 0 TCA(enable) : NO
BIP[10] : 0 Threshold : 0 TCA(enable) : NO
BIP[11] : 0 Threshold : 0 TCA(enable) : NO
BIP[12] : 0 Threshold : 0 TCA(enable) : NO
BIP[13] : 0 Threshold : 0 TCA(enable) : NO
BIP[14] : 0 Threshold : 0 TCA(enable) : NO
BIP[15] : 0 Threshold : 0 TCA(enable) : NO
BIP[16] : 0 Threshold : 0 TCA(enable) : NO
BIP[17] : 0 Threshold : 0 TCA(enable) : NO
BIP[18] : 0 Threshold : 0 TCA(enable) : NO
BIP[19] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[00] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[01] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[02] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[03] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[04] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[05] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[06] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[07] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[08] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[09] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[10] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[11] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[12] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[13] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[14] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[15] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[16] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[17] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[18] : 0 Threshold : 0 TCA(enable) : NO
FRM-ERR[19] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[00] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[01] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[02] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[03] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[04] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[05] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[06] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[07] : 0 Threshold : 0 TCA(enable) : NO

```

```

BAD-SH[08] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[09] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[10] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[11] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[12] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[13] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[14] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[15] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[16] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[17] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[18] : 0 Threshold : 0 TCA(enable) : NO
BAD-SH[19] : 0 Threshold : 0 TCA(enable) : NO
ES : 0 Threshold : 0 TCA(enable) : NO
SES : 0 Threshold : 0 TCA(enable) : NO
UAS : 0 Threshold : 0 TCA(enable) : NO
ES-FE : 0 Threshold : 0 TCA(enable) : NO
SES-FE : 0 Threshold : 0 TCA(enable) : NO
UAS-FE : 0 Threshold : 0 TCA(enable) : NO

```

Last clearing of "show controllers ETHERNET " counters never
RP/0/RP0/CPU0:BH1_P2A4#

This sample displays the history PCS performance monitoring parameters of the 100GE controller at 30-second intervals.

```

RP/0/RP0/CPU0:ios#show controllers hundredGigEctrlr 0/0/0/2 pm history 30-sec pcs 1
Tue Nov 19 09:27:49.169 UTC

```

Ethernet PCS in the current interval [09:27:00 - 09:27:30 Tue Nov 19 2019]

Ethernet PCS current bucket type : Valid

```

BIP[00] : 0
BIP[01] : 0
BIP[02] : 0
BIP[03] : 0
BIP[04] : 0
BIP[05] : 0
.....
BIP[16] : 0
BIP[17] : 0
BIP[18] : 0
BIP[19] : 0
FRM-ERR[00] : 0
FRM-ERR[01] : 0
FRM-ERR[02] : 0
..... 0
FRM-ERR[15] : 0
FRM-ERR[16] : 0
FRM-ERR[17] : 0
FRM-ERR[18] : 0
FRM-ERR[19] : 0
BAD-SH[00] : 0
BAD-SH[01] : 0
BAD-SH[02] : 0
.....
BAD-SH[18] : 0
BAD-SH[19] : 0
ES : 0
SES : 0
UAS : 0
ES-FE : 0
SES-FE : 0
UAS-FE : 0

```

Last clearing of "show controllers ETHERNET " counters never
 RP/0/RP0/CPU0:BH1_P2A4#

This sample displays the current performance monitoring parameters of the optics controller at 10-second intervals as flexi-bin.

RP/0/RP0/CPU0:ios#show controllers optics 0/0/0/0 pm current flex-bin optics 1
 Thu May 21 07:43:38.964 UTC

Optics in the current interval [07:43:30 - 07:43:38 Thu May 21 2020]

Flexible bin interval size: 10 seconds

Optics current bucket type : Valid

	MIN	AVG	MAX	Operational	Configured	TCA	Operational
	Configured	TCA		Threshold(min)	Threshold(min)	(min)	Threshold(max)
	Threshold(max) (max)						
LBC[%]	: 0.0	0.0	0.0	0.0	NA	NO	0.0
	NA	NO					
OPT[dBm]	: -0.13	-0.10	-0.06	0.00	NA	NO	0.00
	NA	NO					
OPR[dBm]	: -3.01	-2.96	-2.92	0.00	NA	NO	0.00
	NA	NO					
CD[ps/nm]	: -3	-2	-1	0	NA	NO	0
	NA	NO					
DGD[ps]	: 1.00	1.67	2.00	0.00	NA	NO	0.00
	NA	NO					
SOPMD[ps^2]	: 17.00	37.00	81.00	0.00	NA	NO	0.00
	NA	NO					
OSNR[dB]	: 37.60	37.60	37.60	0.00	NA	NO	0.00
	NA	NO					
PDL[dB]	: 0.60	0.66	0.70	0.00	NA	NO	0.00
	NA	NO					
PCR[rad/s]	: 0.00	29.11	80.00	0.00	NA	NO	0.00
	NA	NO					
RX_SIG[dBm]	: -3.49	-3.41	-3.36	0.00	NA	NO	0.00
	NA	NO					
FREQ_OFF[Mhz]	: 191	241	301	0	NA	NO	0
	NA	NO					
SNR[dB]	: 14.50	14.62	14.70	0.00	NA	NO	0.00
	NA	NO					
SNR-AX[dB]	: 17.10	17.19	17.30	0.00	NA	NO	0.00
	NA	NO					
SNR-AY[dB]	: 11.90	12.06	12.10	0.00	NA	NO	0.00
	NA	NO					
SNR-BX[dB]	: 0.00	0.00	0.00	0.00	NA	NO	0.00
	NA	NO					
SNR-BY[dB]	: 0.00	0.00	0.00	0.00	NA	NO	0.00
	NA	NO					
SOP-S1	: 0.50	0.55	0.59	0.00	NA	NO	0.00
	NA	NO					
SOP-S2	: -0.59	-0.52	-0.48	0.00	NA	NO	0.00
	NA	NO					
SOP-S3	: -0.67	-0.64	-0.60	0.00	NA	NO	0.00
	NA	NO					

Last clearing of "show controllers OPTICS" counters never

This sample displays the history performance monitoring parameters of the optics controller at 10-second intervals as flexi-bin.

RP/0/RP0/CPU0:ios#show controllers optics 0/0/0/0 pm history flex-bin optics 1 bucket 1

Thu May 21 07:45:44.358 UTC

Optics in interval 1 [07:45:30 - 07:45:40 Thu May 21 2020]

Flexible bin interval size: 10 seconds

Optics history bucket type : Valid

	MIN	AVG	MAX
LBC[%]	: 0.0	0.0	0.0
OPT[dBm]	: -0.12	-0.10	-0.04
OPR[dBm]	: -3.01	-2.97	-2.91
CD[ps/nm]	: -5	-4	-3
DGD[ps]	: 1.00	1.50	2.00
SOPMD[ps^2]	: 28.00	43.10	66.00
OSNR[dB]	: 37.60	37.60	37.60
PDL[dB]	: 0.60	0.65	0.70
PCR[rad/s]	: 0.00	25.70	75.00
RX_SIG[dBm]	: -3.49	-3.44	-3.37
FREQ_OFF[Mhz]	: 235	272	330
SNR[dB]	: 14.60	14.64	14.80
SNR-AX[dB]	: 17.20	17.25	17.30
SNR-AY[dB]	: 11.90	12.02	12.20
SNR-BX[dB]	: 0.00	0.00	0.00
SNR-BY[dB]	: 0.00	0.00	0.00
SOP-S1	: 0.50	0.53	0.57
SOP-S2	: -0.58	-0.53	-0.49
SOP-S3	: -0.69	-0.65	-0.61

Viewing PM Statistics

To view PM statistics for the Ethernet controllers, use this command:

```
RP/0/RP0/CPU0:ios#show controllers HundredGigECtrlr 0/0/0/2 stats
Fri Aug 30 13:10:33.123 IST
Statistics for interface HundredGigECtrlr0/0/0/2 (cached values):
```

```
Ingress:
  Input total bytes          = 1702197139760640
  Input good bytes          = 1702197139760640

  Input total packets       = 13298415154380
  Input 802.1Q frames       = 0
  Input pause frames        = 0
  Input pkts 64 bytes       = 0
  Input pkts 65-127 bytes   = 0
  Input pkts 128-255 bytes  = 13298415154380
  Input pkts 256-511 bytes  = 0
  Input pkts 512-1023 bytes = 0
  Input pkts 1024-1518 bytes = 0
  Input pkts 1519-Max bytes = 0

  Input good pkts           = 13298415154380
  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 encaps = 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             = 1702197139760640
Output good bytes              = 1702197139760640

Output total packets           = 13298415154380
Output 802.1Q frames          = 0
Output pause frames           = 0
Output pkts 64 bytes           = 0
Output pkts 65-127 bytes      = 0
Output pkts 128-255 bytes     = 13298415154380
Output pkts 256-511 bytes     = 0
Output pkts 512-1023 bytes    = 0
Output pkts 1024-1518 bytes   = 0
Output pkts 1519-Max bytes    = 0

Output good pkts               = 13298415154380
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

RP/0/RP0/CPU0:ios#

```



Note Performance monitoring statistics are not supported for the input unicast packets, output unicast packets, and input error fragments counters for Ethernet clients.

Clear PM parameters

Use this task to restore PM parameters on Ethernet or Coherent DSP controllers, ensuring that performance metrics are reset for accurate monitoring.

Clearing PM parameters is necessary when you want to reset the accumulated performance statistics for your controllers. This allows you to start fresh measurements and can be useful for troubleshooting, maintenance, or validating recent changes.

Procedure

Use the command **clear controller *controllertype R/S/I/P* pm**, to clear the performance monitoring parameters for Ethernet and Coherent DSP controllers.

Example:

This sample clears the PM parameters on the Coherent DSP controller.

```
RP/0/RP0/CPU0:ios#show controller CD 0/0/0/0 pm current 15-min fec
Mon Jun 10 11:43:39.981 UTC
```

```
g709 FEC in the current interval [11:30:00 - 11:43:40 Mon Jun 10 2019]
```

```
FEC current bucket type : Invalid
  EC-BITS      : 308360273          Threshold : 903330          TCA(enable) :
YES
  UC-WORDS    : 131108352          Threshold : 5                TCA(enable) :
YES
```

	MIN	AVG	MAX	Threshold (min)	TCA (enable)	Threshold (max)	TCA (enable)
PreFEC BER	3.44E-02	3.45E-02	3.45E-02	0E-15	NO	0E-15	NO
PostFEC BER	0E-15	0E-15	0E-15	0E-15	NO	0E-15	NO
Q	0.51	0.51	0.51	0.00	NO	0.00	NO
Q_Margin	0.00	0.00	0.00	0.00	NO	0.00	NO

```
Last clearing of "show controllers OTU" counters never
```

```
RP/0/RP0/CPU0:ios#clear controller coherentDSP 0/0/0/0 pm
Mon Jun 10 11:44:31.650 UTC
```

```
RP/0/RP0/CPU0:ios#show controller CD 0/0/0/0 pm current 15-min fec
Mon Jun 10 11:44:38.804 UTC
```

```
g709 FEC in the current interval [11:30:00 - 11:44:38 Mon Jun 10 2019]
```

```
FEC current bucket type : Invalid
  EC-BITS      : 0                Threshold : 903330          TCA(enable) :
YES
  UC-WORDS    : 0                Threshold : 5                TCA(enable) :
YES
```

	MIN	AVG	MAX	Threshold (min)	TCA (enable)	Threshold (max)	TCA (enable)
PreFEC BER	3.44E-02	3.44E-02	3.45E-02	0E-15	NO	0E-15	NO
PostFEC BER	0E-15	0E-15	0E-15	0E-15	NO	0E-15	NO
Q	0.51	0.51	0.51	0.00	NO	0.00	NO
Q_Margin	0.00	0.00	0.00	0.00	NO	0.00	NO

```
Last clearing of "show controllers OTU" counters 00:00:07
```

This sample clears the PM parameters on the Ethernet controller.

```
RP/0/RP0/CPU0:ios#clear controller HundredGigECtrlr 0/0/0/2 pm
```



CHAPTER 6

IP Access Lists

- [IP access control lists, on page 91](#)

IP access control lists

An IP Access Control List (ACL) is a network traffic filtering mechanism that

- consists of ordered permit and deny statements evaluating IP addresses, protocols, and directions,
- controls packets by allowing or blocking their movement based on matching criteria,
- and is applied to interfaces, affecting traffic entering or leaving a system, but not traffic originating from the system itself.

ACLs are sequential lists that contain permit or deny statements for IP addresses and upper-layer protocols. These lists control packet flow by matching access list parameters to the information in the packet header. An access list becomes effective only when it is created and applied to an interface.

Packets can be filtered when they arrive at the device (ingress) or leave the device (egress). However, access lists cannot control traffic that originates from the system.

ACL processing paths:

There are two paths for interface packet filtering:

- Hardware programming path (fast path): Uses Ternary Content Addressable Memory (TCAM) via the packet filter Execution Agent for rapid ACL processing.
- Software programming path (slow path): Requires additional configuration through Interface Manager and NetIO, used for management interfaces.

Statistics for ACLs are collected separately for fast path and slow path packets. ACL information is stored globally on the route processor.

Examples of supported features in Cisco NCS 1004:

- Ingress ACL is supported for both IPv4 and IPv6.
- Management interface uses the slow packet path for ACL processing.
- Egress ACL: Self-originated packets are not controlled by ACLs, as these are already managed by the user. ACLs only filter forwarded packets/traffic, for both IPv4 and IPv6.

Best practices and requirements for configuring IP access control lists

These best practices, requirements, and guidelines should be followed when configuring IP ACLs.

- Requirement: Always include at least one permit statement in an access list; otherwise, all packets are denied.
- Best practice: Pay careful attention to the order of permit and deny statements. The order is critical because different orders may lead to different packet outcomes.
- Requirement: Only one access list is allowed for each interface, protocol, and direction.
- Best practice: Apply access lists only to management interfaces. ACLs do not filter traffic on other interfaces or controllers.
- Best practice: Always create the access list before attempting to apply it to an interface using the access-group command.
- Caution: To remove an access list, first remove its reference from the access group, and then remove the access list itself.
- Note: IP access lists cannot filter packets originating from the device itself; they filter only packets arriving at or leaving through an interface.

When you follow these best practices, IP access lists function as intended and protect network resources, preventing unintended packet loss and security breaches.

Configure an IP ACL

Use this procedure whenever you need to restrict or allow specific traffic types entering or leaving a network interface by specifying detailed access list rules.

Procedure

-
- Step 1** Use the command **interface** *interface-type Rack/Slot/Instance/Port* to configure the interface in global configuration mode.
- Example:**
- ```
RP/0/RP0/CPU0:ios#configure
Fri Oct 20 05:25:58.785 UTC
RP/0/RP0/CPU0:ios(config)#interface mgmtEth 0/RP0/CPU0/1t
```
- Step 2** Use the command **ipv4 address** or **ipv6 address** in the interface configuration mode to configure IPv4 or IPv6 address for the interface.
- Example:**
- ```
RP/0/RP0/CPU0:ios(config-if)#ipv4 address 10.1.1.1 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#ipv6 address 1000::1/64
```
- Step 3** Use the **ipv4 | ipv6 access-group access-list-name {ingress | egress}** command, to configure the ACL at the IPv4 or IPv6 interface in the interface configuration mode.
- Example:**

```
RP/0/RP0/CPU0:ios(config-if)#ipv4 access-group IPV4_ICMP_DENY ingress
RP/0/RP0/CPU0:ios(config-if)#ipv4 access-group IPV4_ROUTER_FWD_TELNET_TRAFFIC_DENY egress
RP/0/RP0/CPU0:ios(config-if)#ipv6 access-group IPV6_SSH_DENY ingress
RP/0/RP0/CPU0:ios(config-if)#ipv6 access-group IPV6_ROUTER_FWD_TELNET_TRAFFIC_DENY egress
```

Step 4 Commit the configuration.

Example:

```
RP/0/RP0/CPU0:ios(config-ipv4-acl)# commit
RP/0/RP0/CPU0:ios(config-ipv6-acl)# commit
```

The IP access list is successfully applied to the selected interface. Traffic entering or leaving the interface is now filtered according to the rules you have defined for IPv4 and IPv6, helping enforce your network security policy.

Verify access control lists

Monitoring the number of packets processed by access control lists (ACLs) helps maintain network security and troubleshoot traffic flow. Use this task to review ACL statistics for both IPv4 and IPv6 to ensure filters are working as expected.

Procedure

Use the **show access-lists ipv4** or **show access-lists ipv6** command to verify the IPv4 or IPv6 ACLs.

Example:

IPv4

```
RP/0/RP0/CPU0:ios#show access-lists ipv4
Wed Jan 17 09:52:12.448 IST
ipv4 access-list IPV4_ICMP_DENY
10 deny icmp any any (8 matches)
20 permit ipv4 any any (106 matches)
ipv4 access-list IPV4_ROUTER_FWD_TELNET_TRAFFIC_DENY
10 deny tcp any any eq telnet (3 matches)
20 permit ipv4 any any (6 matches)
```

Example:

IPv6

```
RP/0/RP0/CPU0:ios#show access-lists ipv6
Wed Jan 17 09:52:14.591 IST
ipv6 access-list IPV6_ROUTER_FWD_TELNET_TRAFFIC_DENY
10 deny tcp any any eq telnet (3 matches)
20 permit ipv6 any any (5 matches)
ipv6 access-list IPV6_SSH_DENY
10 deny tcp any any eq ssh (9 matches)
20 permit ipv6 any any (100 matches)
```



CHAPTER 7

Layer 1 Encryption

This chapter describes how to configure the IKEv2 protocol and layer 1 encryption for NCS 1004.



Note In this chapter, "layer 1 encryption" is referred to as "OTNSec".

Table 8: Feature History

Feature Name	Release Information	Feature Description
Encryption Support on 1.2TL Card	Cisco IOS XR Release 7.3.1	AES 256 GCM authenticated OTNSec encryption on 1.2TL line cards is supported. It uses only pre-shared keys for authentication. Optical encryption secures the communications link in and out of a facility, rendering all data undecipherable to hackers who tap into networks.

The Need for High Speed Encryption

Most of the emphasis on protecting networks today is focused on protecting data within data center. However, the infrastructure of networks that connect these data centers are as vulnerable to calculated attacks as the data centers themselves. As more sensitive information gets transmitted across fiber-optic networks, cyber criminals are increasingly turning their attention to intercepting the data when it travels across the network.

With the increase in network or fiber optic hacks, the need for data protection is paramount. Encryption of any data that leaves the data centers is becoming an important requirement for cloud operators. Optical encryption secures everything on the communications link in and out of a facility rendering all data undecipherable to any hacker that taps into the fiber strand. *Protecting data at high speeds or lines rates is a requirement for data centers today.*


The Cisco NCS 1004 brings to you AES256 based OTNSec encryption for 100GE and OTU4 clients. Encryption is supported on the 1.2T cards.

OTNSec encryption uses the IKEv2 protocol to negotiate and establish the IKEv2 and OTNSec Security Associations (SA). IKEv2 is used for authentication of the devices in an encryption session, and the protocol

provides pre-shared keys (PSK) or RSA certificate-based authentication. The IKEv2 datagrams are carried as payloads using the point-to-point protocol (PPP) over the GCC channel.

To implement this, an IKE session is established between the two endpoints, Site A and Site B, for overhead control plane communication between the two data centers. Data is then encrypted at Site A using OTNSec encryption and decrypted at Site B.

Figure 2: OTNSec Site-to-Site Example and Components




The recommended deployment is to have a single IKEv2 session running over a GCC2 channel per trunk port which creates the child SAs for each of the OTNSec controllers that are configured on the trunk port.

- [IKEv2 Overview, on page 96](#)
- [OTNSec Encryption Overview, on page 98](#)
- [Prerequisites, on page 99](#)
- [Limitations, on page 100](#)
- [Configuration Workflow, on page 100](#)
- [Configuration Example, on page 107](#)
- [Verification, on page 110](#)
- [Troubleshooting, on page 110](#)
- [IKEv2 Certificate-Based Authentication, on page 111](#)
- [You May Be Interested In, on page 115](#)

IKEv2 Overview

Internet Key Exchange Version 2 (IKEv2) is a request and response encryption that establishes and handles security associations (SA) in an authentication suite, such as OTNSec, to ensure secure traffic. IKE performs mutual authentication between two endpoints and establishes an IKE Security Association (SA). All IKE communications consist of pairs of messages that include a request and a response. The pair is called an exchange or a request-response pair. The first two exchanges of messages establishing an IKE SA are called the IKE_SA_INIT exchange and the IKE_AUTH exchange; subsequent IKE exchanges are called either CREATE_CHILD_SA exchanges or INFORMATIONAL exchanges. IKEv2 uses sequence numbers and acknowledgments to provide reliability, and mandates some error-processing logistics and shared state management (windowing). IKEv2 does not process a request until it determines the requester. This helps to mitigate DoS attacks. IKEv2 provides built-in support for Dead Peer Detection (DPD), which periodically confirms the availability of the peer node. When there is no response from the peer node, the system attempts to establish the session again.

Figure 3: IKEv2 Exchanges



IKEv2 is defined in RFC 7296 and consists of the following constructs:

- **Keyring**

A keyring is a repository of symmetric and asymmetric pre-shared keys that is configured for a peer and identified using the IP address of the peer. The keyring is associated with an IKEv2 profile and therefore, caters to a set of peers that match the IKEv2 profile. This is a required configuration for the pre-shared keys authentication method that is used for NCS 1004.



Note The certificate-based authentication that uses RSA signatures can be used instead of the keyring. If both methods of authentication are configured, the certificate-based authentication takes precedence. See [IKEv2 Certificate-Based Authentication, on page 111](#).

- **IKEv2 Profile**

An IKEv2 profile is a repository of nonnegotiable parameters of the IKE SA, such as authentication method and services that are available to the authenticated peers that match the profile. The profile match lookup is done based on the IP address of the remote identity. For security purposes, the IKE SAs have a lifetime that is defined in the IKEv2 profile. The lifetime range, in seconds, is from 120 to 86400. The SAs are rekeyed proactively before the expiry of the lifetime. The default lifetime is 86400. An IKEv2 profile must be attached to an OTNSec configuration on the ODU4 controllers on both the IKEv2 initiator and responder. This is a required configuration.



Note Only one authentication method is supported for the local peer but multiple authentication methods can be configured for the remote peer.

If both methods of authentication are configured, keyring and certificate trustpoint (see, [IKEv2 Certificate-Based Authentication, on page 111](#)) in the profile, the remote peer can authenticate itself using either method. **authentication remote [pre-shared] rsa-signature** can be used to exclusively control the remote authentication method. Similarly, **authentication local [pre-shared] rsa-signature** can be used to exclusively configure local authentication method. If it is not configured, the certificate-based authentication takes precedence.

• IKEv2 Proposal

An IKEv2 proposal is a collection of transforms that are used in the negotiation of IKE SAs as part of the IKE_SA_INIT exchange. The IKE2 proposal must be attached to an IKEv2 policy. This is an optional configuration. The transform types used in the negotiation are as follows:

- Encryption algorithm
- Integrity algorithm
- Pseudo-Random Function (PRF) algorithm
- Diffie-Hellman (DH) group



Note The IKEv2 proposal must have at least one algorithm of each type. It is possible to specify multiple algorithms for each type; the order in which the algorithms are specified determines the precedence.

• IKEv2 Policy

IKEv2 employs policies that are configured on each peer to negotiate handshakes between the two peers. An IKEv2 policy contains proposals that are used to negotiate the encryption, integrity, PRF algorithms, and DH group in the SA_INIT exchange. An IKEv2 policy is selected based on the local IP address. This is an optional configuration.



Note The default IKEv2 proposal is used with default IKEv2 policy in the absence of any user-defined policy.

OTNSec Encryption Overview

OTNSec encryption in NCS 1004 has the following characteristics:

- The OTN layer 1 security is supported over the OPU client payload.

- The Galois-Counter-Mode (GCM) AES 256-bit security is the default cipher used for encryption and decryption of the OPU payloads.
- Each client offers an independent encrypted channel in each direction.
- There are two banks of 256-bit programmable key registers (current key and future key) that permit key updates through the software without interrupting traffic.
Each key is associated with an Association Number [AN(1:0)] allowing up to four different numbers.
- Interhost key exchange is supported through communication over GCC.
- The encryption is supported in headless mode.

The OTNSec control plane generates two different keys, one for the transmit (Tx) side and the other for the receive (Rx) side. These keys are used by the line card to program the encryptor and decryptor blocks. These blocks encrypt and decrypt the data packets between the trunk ports of the two nodes. For security purposes, the keys have a lifetime. A key's lifetime specifies the time the key expires.

The key lifetime for the child SAs can be configured using the sak-rekey-interval which ranges from 30 seconds to 14 days. For example, if the sak-rekey-interval is configured for five minutes, a new key is generated by the OTNSec layer every five minutes. In the absence of a lifetime configuration, the default lifetime is 14.18 days. When the key reaches the maximum lifetime, it becomes invalid and the CRYPTO-KEY-EXPIRED alarm is raised. Volume-based rekeying is supported; it prevents the key from reaching the maximum lifetime. This allows the OTNSec layer to generate a new key when 70% of the lifetime (11 days) of the current key is over.

When the lifetime of the first key expires, it automatically rolls over to the next key. To achieve a hitless rollover, the lifetimes of the keys need to be overlapped so that for a certain period of time both keys are active. To maintain this seamless switchover, a key index table is maintained. Each key pair (Tx and Rx) is associated with an Association Number (AN). The index table allows up to four numbers (0,1, 2, and 3). When the keys are installed, the Rx AN number of node A must match the Tx AN number of node B. Also, the Tx AN number of node A must match the Rx AN number of node B. If there is a mismatch of the AN numbers between the peer nodes, the CRYPTO-INDEX-MISMATCH alarm is raised.

Prerequisites

- Ensure that the required k9sec.rpm package is installed.
- Configure the line card in the muxponder or muxponder slice mode using the following commands:

- **1.2T Card:**

- muxponder mode:

hw-module location *location* mxponder client-rate 100GE | OTU4

hw-module location *location* mxponder trunk-rate {100G | 200G | 300G | 400G | 500G | 600G}

- muxponder slice mode:

hw-module location *location* mxponder-slice *mxponder-slice-number* client-rate 100GE | OTU4

hw-module location *location* mxponder-slice trunk-rate { 100G | 200G | 300G | 400G | 500G | 600G }

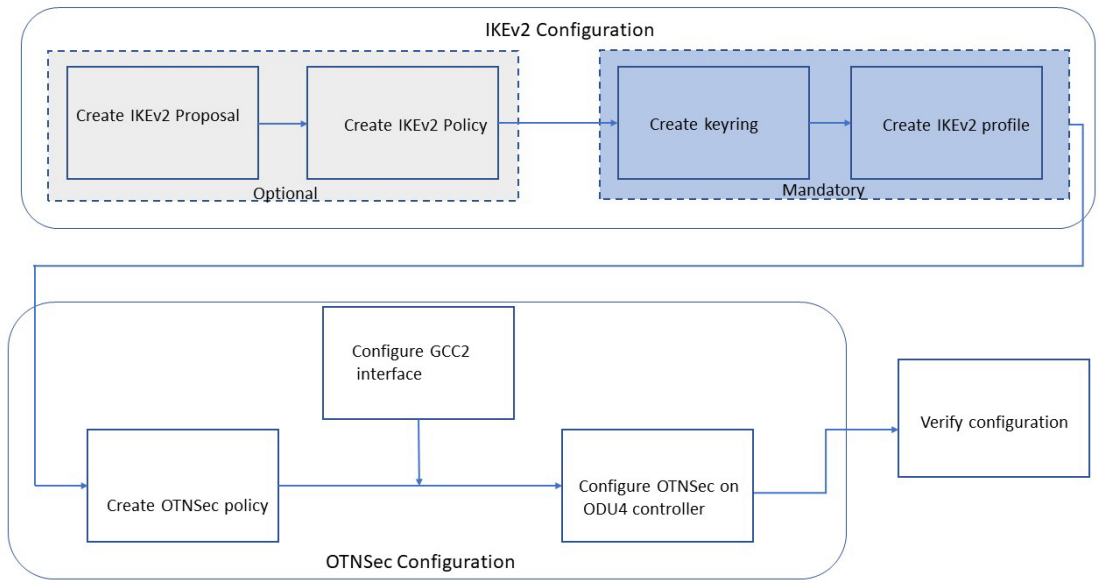
Limitations

- Traffic is impacted for a few seconds if the RP fails or GCC2 control plane goes down, during a key rollover.
- The sak-rekey-interval must be configured on the initiator and responder node.

Configuration Workflow

This section describes the workflow to configure IKEv2 and OTNSec encryption on NCS 1004. The authentication method used is pre-shared keys (PSKs).

Figure 4: L1 Encryption Workflow



369571

Table 9: Workflow for Configuring IKEv2 and OTNSec Encryption on NCS 1004

Workflow Sequence	Details
IKE Configuration	
Configuring an IKEv2 Proposal, on page 101	<p>(Optional) Configure an IKEv2 proposal manually; otherwise, the default IKEv2 proposal is used in the default IKEv2 policy.</p> <p>The default IKEv2 proposal requires no configuration and is a collection of commonly used transforms types, which are as follows:</p> <pre> encryption cbc-aes-256 integrity sha512, sha384 prf sha512, sha384 dh 19, 20, 21 </pre>

Workflow Sequence	Details
Configuring an IKEv2 Policy, on page 102	(Optional) Configure an IKEv2 policy manually; otherwise, the default proposal associated with the default policy is used for negotiation. Note An IKEv2 policy with no proposal is considered incomplete.
Configuring a Keyring, on page 103	Configure a keyring as the local or remote authentication method is a preshared key.
Configuring a IKEv2 Profile, on page 104	Configure an IKEv2 profile. Note <ul style="list-style-type: none"> • The IKEv2 profile must be attached to the OTNSec profile on both the IKEv2 initiator and the responder. • The DPD interval is 10 seconds. If there is no response from the peer node, it retries every two seconds with a maximum of five attempts. After five retries, the IKE session is brought down. NCS 1004 supports headless mode. Therefore, even though the control plane is down, traffic is not impacted because the encryption and decryption keys are still active on the line cards. The data path functions in a locally secure mode and the OTNSEC-LOCALLY-SECURED alarm is raised.
OTNSec Configuration	
Configuring an OTNSec Policy, on page 104	(Optional) Configure the OTNSec policy.
Configuring the GCC Interface, on page 105	Configure the GCC2 interface.
Configuring OTNSec on ODU4 Controllers, on page 106	Configure the ODU4 controller that is mapped to the HundredGigE controller .
Verification	
Verification, on page 110	Verify the IKEv2 and OTNSec configuration.

Configuring an IKEv2 Proposal

To configure an IKEv2 proposal, use the following commands:

```
config
```

```
ikev2 proposal proposal-name
```

```
encryption {aes-gcm-256} {aes-gcm-128} {aes-cbc-256} {aes-cbc-192} {aes-cbc-128}
```

```
integrity {sha-1} {sha-256} {sha-384} {sha-512}
```

```
prf {sha-1} {sha-256} {sha-384} {sha-512}
dh {19} {20} {21}
```



Note Configuring an AES-GCM encryption algorithm does not require configuring an integrity algorithm. AES-GCM and non-GCM algorithms cannot be configured in the same proposal. However, you can configure the AES-GCM and non-GCM algorithms under two different proposals and attach both the proposals to the same IKEv2 policy.

The following sample displays how to configure an IKEv2 proposal.

```
RP/0/RP0/CPU0:ios#configure
Thu Mar  7 19:19:30.259 UTC
RP/0/RP0/CPU0:ios(config)#ikev2 proposal proposal1
RP/0/RP0/CPU0:ios(config-ikev2-proposal-proposal1)#encryption aes-cbc-256
RP/0/RP0/CPU0:ios(config-ikev2-proposal-proposal1)#integrity sha-1
RP/0/RP0/CPU0:ios(config-ikev2-proposal-proposal1)#prf sha-256
RP/0/RP0/CPU0:ios(config-ikev2-proposal-proposal1)#dh 20
RP/0/RP0/CPU0:ios(config-ikev2-proposal-proposal1)#commit
Thu Mar  7 19:20:30.916 UTC
RP/0/RP0/CPU0:ios(config-ikev2-proposal-proposal1)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show ikev2 proposal proposal1
Thu Mar  7 19:20:48.929 UTC

Proposal Name           : proposal1
=====
Status                  : Complete
-----
Total Number of Enc. Alg. : 1
  Encr. Alg.             : CBC-AES-256
-----
Total Number of Hash. Alg. : 1
  Hash. Alg.             : SHA 1
-----
Total Number of PRF. Alg. : 1
  PRF. Alg.              : SHA 256
-----
Total Number of DH Group : 1
  DH Group                : Group 20
```

Configuring an IKEv2 Policy

To configure an IKEv2 policy, use the following commands:

config

ikev2 policy *policy-name*

proposal *proposal-name1 proposal-name2 proposal-name3*

match address local { *ipv4-address* }

The following sample displays how to configure an IKEv2 policy.

```
RP/0/RP0/CPU0:ios#configure
Thu Mar  7 19:26:45.752 UTC
RP/0/RP0/CPU0:ios(config)#ikev2 policy mypolicy
RP/0/RP0/CPU0:ios(config-ikev2-policy-mypolicy)#proposal proposal1
RP/0/RP0/CPU0:ios(config-ikev2-policy-mypolicy)#match address local 10.1.1.1
```

```
RP/0/RP0/CPU0:ios(config-ikev2-policy-mypolicy)#commit
Thu Mar  7 19:29:25.043 UTC
RP/0/RP0/CPU0:ios(config-ikev2-policy-mypolicy)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show ikev2 policy mypolicy
Thu Mar  7 19:30:30.343 UTC
```

```
Policy Name                               : mypolicy
=====
Total number of match local addr          : 1
  Match address local                     : 10.1.1.1
-----
Total number of proposal attached         : 1
  Proposal Name                           : proposal1
```

Configuring a Keyring

To configure a keyring, use the following commands:

config

keyring *keyring-name*

peer *peer-block name*

address *{ipv4-address [mask]}*



Note The IP address of the far-end node (remote node) must be used.

pre-shared-key *{{key} {clear clear-text key} {local local key} {passwordencrypted key}}*



Note The key input can either be in clear text or in type 7 encrypted password format.

The following sample displays how to configure a keyring.

```
RP/0/RP0/CPU0:ios#configure
Thu Mar  7 19:33:14.594 UTC
RP/0/RP0/CPU0:ios(config)#keyring kyr1
RP/0/RP0/CPU0:ios(config-keyring-kyr1)#peer peer1
RP/0/RP0/CPU0:ios(config-keyring-kyr1-peer-peer1)#address 10.1.1.2 255.255.255.0
RP/0/RP0/CPU0:ios(config-keyring-kyr1-peer-peer1)#pre-shared-key password 106D000A064743595F
RP/0/RP0/CPU0:ios(config-keyring-kyr1-peer-peer1)#commit
Thu Mar  7 19:54:33.314 UTC
RP/0/RP0/CPU0:ios(config-keyring-kyr1-peer-peer1)#exit
RP/0/RP0/CPU0:ios(config-keyring-kyr1)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show keyring kyr1
Thu Mar  7 19:58:07.135 UTC
```

```
Keyring Name                               : kyr1
=====
Total Peers                               : 1
-----
Peer Name                                 : peer1
IP Address                                : 10.1.1.2
Subnet Mask                               : 255.255.255.0
Identity                                  : Not configured
```

```

Local PSK                               : Configured
Remote PSK                               : Configured
PPK Mode                                 : Manual
PPK Mandatory                             : Yes
Local PSK Hash (Algo:SHA256, Format:base64) :
D8XsPtNaX4gCSp4bCNLqHkgEWMJs9l6v2aXirCvGK2I=
Remote PSK Hash (Algo:SHA256, Format:base64) :
jBFwTDBOK/kT894SrK3T8hJQ5BZtCus/KyEgNj4cJVQ=
Manual PPK Hash (Algo:SHA256, Format:base64) :
6c7nGrky/ehjM40Ivk3p3+OeoEm9r7NCzmWexUULaa4=

```

Configuring a IKEv2 Profile

To configure an IKEv2 profile, use the following commands:

config

ikev2 profile *profile-name*

match identity remote address *{ipv4-address [mask]}*

keyring *keyring-name*

lifetime *seconds*



Note The lifetime range, in seconds, is from 120 to 86400.

The following sample displays how to configure an IKEv2 profile.

```

RP/0/RP0/CPU0:ios#configure
Thu Mar  7 20:00:36.490 UTC
RP/0/RP0/CPU0:ios(config)#ikev2 profile profile1
RP/0/RP0/CPU0:ios(config-ikev2-profile-profile1)#match identity remote address 10.1.1.2
255.255.255.0
RP/0/RP0/CPU0:ios(config-ikev2-profile-profile1)#keyring kyr1
RP/0/RP0/CPU0:ios(config-ikev2-profile-profile1)#lifetime 86400
RP/0/RP0/CPU0:ios(config-ikev2-profile-profile1)#commit
Thu Mar  7 20:15:03.401 UTC
RP/0/RP0/CPU0:ios(config-ikev2-profile-profile1)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show ikev2 profile profile1
Thu Mar  7 20:15:25.776 UTC

```

```

Profile Name                               : profile1
=====
Keyring                                    : kyr1
Lifetime (Sec)                             : 120
DPD Interval (Sec)                         : 10
DPD Retry Interval (Sec)                   : 2
Match ANY                                   : NO
Total Match remote peers                   : 1
  Addr/Prefix                               : 10.1.1.2/255.255.255.0

```

Configuring an OTNSec Policy

To configure an OTNSec policy, use the following commands:

config

otnsec-policy *policy-name*
cipher-suite **AES-GCM-256**
security-policy **must-secure**
sak-rekey-interval *seconds*



Note The interval range, in seconds, is from 30 to 1209600. SAK rekey timer does not start by default until it is configured.

The following sample displays how to configure an OTNSec policy.

```
RP/0/RP0/CPU0:ios#configure
Mon Mar 11 15:16:58.417 UTC
RP/0/RP0/CPU0:ios(config)#otnsec policy otnsec-policy1
RP/0/RP0/CPU0:ios(config-otnsec-policy)#cipher-suite AES-GCM-256
RP/0/RP0/CPU0:ios(config-otnsec-policy)#security-policy must-secure
RP/0/RP0/CPU0:ios(config-otnsec-policy)#sak-rekey-interval 120
RP/0/RP0/CPU0:ios(config-otnsec-policy)#commit
```

The following is a sample of an OTNSec policy.

```
RP/0/RP0/CPU0:ios#show run otnsec policy otnsec-policy1
Tue Mar 12 11:14:03.591 UTC
otnsec policy otnsec-policy1
  cipher-suite AES-GCM-256
  security-policy must-secure
  sak-rekey-interval 120
!
```



Note When a software upgrade is performed from R.7.0.1 to later releases, traffic is impacted. This happens if the sak-rekey-interval is configured. To prevent traffic loss, disable the sak-rekey-interval before the software upgrade using the following commands:

```
Tue Nov 26 12:41:01.768 IST
RP/0/RP0/CPU0:ios(config)#otnsec policy OP1
RP/0/RP0/CPU0:ios(config-otnsec-policy)#no sak-rekey-interval
```

The sak-rekey-interval can be configured again after the upgrade process is complete.

Configuring the GCC Interface

To configure the GCC interface, use the following commands:

```
config
interface GCC2 R/S/I/P
ipv4 address ipv4-address
```

The following sample displays how to configure the GCC2 interface.

```
RP/0/RP0/CPU0:ios#config
Tue Mar 12 12:06:32.547 UTC
RP/0/RP0/CPU0:ios(config)#controller odu4 0/1/0/0/1
RP/0/RP0/CPU0:ios(config-odu4)#gcc2
```

```

RP/0/RP0/CPU0:ios(config-odu4)#commit
RP/0/RP0/CPU0:ios(config-odu4)#exit

RP/0/RP0/CPU0:ios#config
Tue Mar 12 11:16:04.749 UTC
RP/0/RP0/CPU0:ios(config)#interface GCC2 0/1/0/0/1
P/0/RP0/CPU0:ios(config-if)#ipv4 address 10.1.1.1 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#commit
Tue Mar 12 11:18:32.867 UTC
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#sh run interface gcc2 0/1/0/0/1
Tue Mar 12 11:19:00.475 UTC
interface GCC20/1/0/0/1
  ipv4 address 10.1.1.1 255.255.255.0
!

RP/0/RP0/CPU0:ios#config
Wed Sep 28 23:10:28.258 UTC
RP/0/RP0/CPU0:ios(config)#controller ODU4 0/0/0/12
RP/0/RP0/CPU0:ios(config-oduc4)#gcc2
RP/0/RP0/CPU0:ios(config-oduc4)#commit
RP/0/RP0/CPU0:ios(config-oduc4)#exit

RP/0/RP0/CPU0:ios#config
Wed Sep 28 23:10:29.808 UTC
RP/0/RP0/CPU0:ios(config)#interface GCC2 0/0/0/12
P/0/RP0/CPU0:ios(config-if)#ipv4 address 10.1.1.1 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#commit
Wed Sep 28 23:10:30.260 UTC UTC
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#sh run interface gcc2 0/0/0/12
Tue Mar 12 11:19:00.475 UTC
interface GCC20/0/0/12
  ipv4 address 10.1.1.1 255.255.255.0
!

```

Configuring OTNSec on ODU4 Controllers

To configure the OTNSec on ODU4 controller, use the following commands:

config

controller ODU4 *rack/slot/instance/port*

otnsec

source ipv4 *ipv4-address*

destination ipv4 *ipv4-address*

session-id *session-id*

policy *policy-name*

ikev2 *profile-name*



Note The session ID ranges 1–65535.

The following sample displays how to configure OTNSec on the ODU4 controller.

```

RP/0/RP0/CPU0:ios#configure
Mon Mar 12 12:10:21.374 UTC
RP/0/RP0/CPU0:ios(config)#controller ODU4 0/1/0/0/1
RP/0/RP0/CPU0:ios(config-odu4)#otnsec
RP/0/RP0/CPU0:ios(config-otnsec)#source ipv4 10.1.1.1
RP/0/RP0/CPU0:ios(config-otnsec)#destination ipv4 10.1.1.2
RP/0/RP0/CPU0:ios(config-otnsec)#session-id 9000
RP/0/RP0/CPU0:ios(config-otnsec)#policy otnsec-policy1
RP/0/RP0/CPU0:ios(config-otnsec)#ikev2 profile1
RP/0/RP0/CPU0:ios(config-otnsec)#commit
Mon Mar 12 12:14:17.609 UTC
RP/0/RP0/CPU0:ios(config-otnsec)#exit
RP/0/RP0/CPU0:ios(config)#exit

```

Configuration Example

In the following example, there are two nodes. The node with the lower IP address always acts as the initiator. In this case, node A (SITE-A) has the role of an initiator while Node B (SITE-B) has the role of a responder. In this example, the default IKE proposal and policy have been used on both nodes.

Figure 5: Configuration Schema



The configuration on Node A is displayed below.

Node A (Initiator)
Keyring
<pre> RP/0/RP0/CPU0:SITE-A#configure RP/0/RP0/CPU0:SITE-A(config)#keyring KR1 RP/0/RP0/CPU0:SITE-A(config-keyring-KR1)#peer SITE-B RP/0/RP0/CPU0:SITE-A(config-keyring-KR1-peer-SITE-B)#address 10.1.1.2 255.255.255.0 RP/0/RP0/CPU0:SITE-A(config-keyring-KR1-peer-SITE-B)#pre-shared-key password 106D000A064743595 RP/0/RP0/CPU0:SITE-A(config-keyring-KR1-peer-SITE-B)#commit RP/0/RP0/CPU0:SITE-A(config-keyring-KR1-peer-SITE-B)#exit RP/0/RP0/CPU0:SITE-A(config-keyring-KR1)#exit </pre>
IKEv2 profile
<pre> RP/0/RP0/CPU0:SITE-A(config)#ikev2 profile IP1 RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#match identity remote address 10.1.1.2 255.255.255.0 RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#keyring KR1 RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#lifetime 86400 RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#commit RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#exit </pre>

Node A (Initiator)
OTNSec policy
<pre>RP/0/RP0/CPU0:SITE-A(config)#otnsec policy OP1 RP/0/RP0/CPU0:SITE-A(config-otnsec-policy)#cipher-suite AES-GCM-256 RP/0/RP0/CPU0:SITE-A(config-otnsec-policy)#security-policy must-secure RP/0/RP0/CPU0:SITE-A(config-otnsec-policy)#sak-rekey-interval 120 RP/0/RP0/CPU0:SITE-A(config-otnsec-policy)#commit RP/0/RP0/CPU0:SITE-A(config-otnsec-policy)#exit</pre>
GCC interface
<pre>RP/0/RP0/CPU0:SITE-A(config)#controller odu4 0/1/0/0/1 RP/0/RP0/CPU0:SITE-A(config-odu4)#gcc2 RP/0/RP0/CPU0:SITE-A(config-odu4)#commit RP/0/RP0/CPU0:SITE-A(config-odu4)#exit RP/0/RP0/CPU0:SITE-A(config)#interface GCC2 0/1/0/0/1 RP/0/RP0/CPU0:SITE-A(config-if)#ipv4 address 10.1.1.1 255.255.255.0 RP/0/RP0/CPU0:SITE-A(config-if)#commit RP/0/RP0/CPU0:SITE-A(config-if)#exit</pre>
OTNSec on ODU4 controller
<pre>RP/0/RP0/CPU0:SITE-A(config)#controller odu4 0/1/0/0/1 RP/0/RP0/CPU0:SITE-A(config-odu4)#otnsec RP/0/RP0/CPU0:SITE-A(config-otnsec)#source ipv4 10.1.1.1 RP/0/RP0/CPU0:SITE-A(config-otnsec)#destination ipv4 10.1.1.2 RP/0/RP0/CPU0:SITE-A(config-otnsec)#session-id 9000 RP/0/RP0/CPU0:SITE-A(config-otnsec)#policy OP1 RP/0/RP0/CPU0:SITE-A(config-otnsec)#ikev2 IP1 RP/0/RP0/CPU0:SITE-A(config-otnsec)#commit RP/0/RP0/CPU0:SITE-A(config-otnsec)#exit RP/0/RP0/CPU0:SITE-A(config-odu4)#exit RP/0/RP0/CPU0:SITE-B(config)#exit</pre>

The configuration on Node B is displayed below.

Node B (Responder)
Keyring
<pre>RP/0/RP0/CPU0:SITE-B#configure RP/0/RP0/CPU0:SITE-B(config)#keyring KR1 RP/0/RP0/CPU0:SITE-B(config-keyring-KR1)#peer SITE-A RP/0/RP0/CPU0:SITE-B(config-keyring-KR1-peer-SITE-A)#address 10.1.1.1 255.255.255.0 RP/0/RP0/CPU0:SITE-B(config-keyring-KR1-peer-SITE-A)#pre-shared-key password 14341B180F547B7977 RP/0/RP0/CPU0:SITE-B(config-keyring-KR1-peer-SITE-A)#commit RP/0/RP0/CPU0:SITE-B(config-keyring-KR1-peer-SITE-A)#exit RP/0/RP0/CPU0:SITE-B(config-keyring-KR1)#exit</pre>
IKEv2 profile

Node B (Responder)

```
RP/0/RP0/CPU0:SITE-B(config)#ikev2 profile IP1
RP/0/RP0/CPU0:SITE-B(config-ikev2-profile-IP1)#match identity remote address 10.1.1.1
255.255.255.0
RP/0/RP0/CPU0:SITE-B(config-ikev2-profile-IP1)#keyring KR1
RP/0/RP0/CPU0:SITE-B(config-ikev2-profile-IP1)#lifetime 86400
RP/0/RP0/CPU0:SITE-B(config-ikev2-profile-IP1)#commit
RP/0/RP0/CPU0:SITE-B(config-ikev2-profile-IP1)#exit
```

OTNSec policy

```
RP/0/RP0/CPU0:SITE-B(config)#otnsec policy OP1
RP/0/RP0/CPU0:SITE-B(config-otnsec-policy)#cipher-suite AES-GCM-256
RP/0/RP0/CPU0:SITE-B(config-otnsec-policy)#security-policy must-secure
RP/0/RP0/CPU0:SITE-B(config-otnsec-policy)#sak-rekey-interval 120
RP/0/RP0/CPU0:SITE-B(config-otnsec-policy)#commit
RP/0/RP0/CPU0:SITE-B(config-otnsec-policy)#exit
```

GCC interface

```
RP/0/RP0/CPU0:SITE-B(config)#controller odu4 0/1/0/0/1
RP/0/RP0/CPU0:SITE-B(config-odu4)#gcc2
RP/0/RP0/CPU0:SITE-B(config-odu4)#commit
RP/0/RP0/CPU0:SITE-B(config-odu4)#exit
RP/0/RP0/CPU0:SITE-B(config)#interface GCC2 0/1/0/0/1
RP/0/RP0/CPU0:SITE-B(config-if)#ipv4 address 10.1.1.2 255.255.255.0
RP/0/RP0/CPU0:SITE-B(config-if)#commit
RP/0/RP0/CPU0:SITE-B(config-if)#exit
```

GCC interface for OTN-XP card

```
RP/0/RP0/CPU0:SITE-B(config)#controller oduc4 0/0/0/12
RP/0/RP0/CPU0:SITE-B(config-oduc4)#gcc2
RP/0/RP0/CPU0:SITE-B(config-oduc4)#commit
RP/0/RP0/CPU0:SITE-B(config-oduc4)#exit
RP/0/RP0/CPU0:SITE-B(config)#interface GCC2 0/0/0/12
RP/0/RP0/CPU0:SITE-B(config-if)#ipv4 address 10.1.1.2 255.255.255.0
RP/0/RP0/CPU0:SITE-B(config-if)#commit
RP/0/RP0/CPU0:SITE-B(config-if)#exit
```

OTNSec on ODU4 controller

```
RP/0/RP0/CPU0:SITE-B(config)#controller odu4 0/1/0/0/1
RP/0/RP0/CPU0:SITE-B(config-odu4)#otnsec
RP/0/RP0/CPU0:SITE-B(config-otnsec)#source ipv4 10.1.1.2
RP/0/RP0/CPU0:SITE-B(config-otnsec)#destination ipv4 10.1.1.1
RP/0/RP0/CPU0:SITE-B(config-otnsec)#session-id 9000
RP/0/RP0/CPU0:SITE-B(config-otnsec)#policy OP1
RP/0/RP0/CPU0:SITE-B(config-otnsec)#ikev2 IP1
RP/0/RP0/CPU0:SITE-B(config-otnsec)#commit
RP/0/RP0/CPU0:SITE-B(config-otnsec)#exit
RP/0/RP0/CPU0:SITE-B(config-odu4)#exit
RP/0/RP0/CPU0:SITE-B(config)#exit
```

OTNSec on ODOC4 controller

Node B (Responder)

```

RP/0/RP0/CPU0:SITE-B(config)#controller oduc4 0/0/0/12
RP/0/RP0/CPU0:SITE-B(config-oduc4)#otnsec
RP/0/RP0/CPU0:SITE-B(config-otnsec)#source ipv4 10.1.1.2
RP/0/RP0/CPU0:SITE-B(config-otnsec)#destination ipv4 10.1.1.1
RP/0/RP0/CPU0:SITE-B(config-otnsec)#session-id 99
RP/0/RP0/CPU0:SITE-B(config-otnsec)#policy OP1
RP/0/RP0/CPU0:SITE-B(config-otnsec)#ikev2 IP1
RP/0/RP0/CPU0:SITE-B(config-otnsec)#commit
RP/0/RP0/CPU0:SITE-B(config-otnsec)#exit
RP/0/RP0/CPU0:SITE-B(config-oduc4)#exit
RP/0/RP0/CPU0:SITE-B(config)#exit

```

Verification

- Verify that there are no alarms on the ports of the NCS 1004.
- Use the **show** commands listed in the table below to verify the IKEv2 and OTNSec configuration. For details of these commands, see the *Command Reference for Cisco NCS 1004*.

Table 10: Show Commands

Show Commands	Purpose
show run ikev2	Displays the running configuration of IKEv2
show ikev2 session	Displays the child SAs created for the session
show ip interface brief	Displays the status of the GCC interfaces
show run controller ODU4 0/1/0/0/1	Displays the running configuration of the ODU4 controller
show controllers ODU4 0/1/0/0/1 otnsec	Displays the OTNSec configuration on the ODU4 controller
show controllers ODU4 0/1/0/0/1 pm current 15-min otnsec	Displays the PM statistics that help verify the encrypted and decrypted blocks.

Troubleshooting

Problem: The IKE session is not established between the two nodes.

Solution: Check the status of the GCC interface using the **show ip interface brief** command.

To gather logs and traces, use the **show tech-support ncs1004 detail**, **show tech-support ikev2**, and **show tech-support otnsec** commands.

IKEv2 Certificate-Based Authentication

IKEv2 can use RSA digital signatures to authenticate peer devices before setting up SAs. RSA signatures employ a PKI-based method of authentication.

Certification Authority (CA) interoperability permits Cisco NCS 1004 devices and CAs to communicate so that your device can obtain and use digital certificates from the CA. A CA is responsible for managing certificate requests and issuing certificates to participating network devices. With a CA, a router authenticates itself to the remote router by sending a certificate to the remote router and performing some public key cryptography. Each router must send its own unique certificate that was issued and validated by the CA. This process works because the certificate of each router encapsulates the public key of the router, each certificate is authenticated by the CA, and all participating routers recognize the CA as an authenticating authority. This scheme is called IKE with an RSA signature.

In public key cryptography, such as the RSA encryption system, each user has a key pair containing both a public and a private key. The keys act as complements, and anything encrypted with one of the keys can be decrypted with the other. In simple terms, a signature is formed when data is encrypted with a user's private key. The receiver verifies the signature by decrypting the message with the sender's public key. The fact that the message could be decrypted using the sender's public key indicates that the holder of the private key, the sender, must have created the message. This process relies on the receiver's having a copy of the sender's public key and knowing with a high degree of certainty that it does belong to the sender and not to someone pretending to be the sender.

Configuring IKEv2 Certificate-Based Authentication

To configure IKEv2 certificate-based authentication, perform the following steps:

1. Configure router hostname and IP domain name—You must configure the hostname and IP domain name of the router if they have not already been configured. The hostname and IP domain name are required because the router assigns a fully qualified domain name (FQDN) to the keys and certificates used by OTNsec, and the FQDN is based on the hostname and IP domain name you assign to the router.

configure

hostname name

domain name domain-name

commit

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:IOS(config)#hostname myhost
RP/0/RP0/CPU0:IOS(config)#domain name mydomain.com
RP/0/RP0/CPU0:IOS(config)#commit
```

2. Generate RSA key pair—The RSA key pair is required before you can obtain a certificate for your router.

crypto key generate rsa keypair-label

```
RP/0/RP0/CPU0:ios#crypto key generate rsa tp
Thu May 7 16:18:44.243 IST
The name for the keys will be: tp
Do you really want to replace them? [yes/no]: yes
Choose the size of the key modulus in the range of 512 to 4096 for your General Purpose Keypair. Choosing a key modulus greater than 512 may take a few minutes.
```

```

How many bits in the [2048]:
Generating RSA keys ...
Done w/ crypto generate keypair
[OK]

RP/0/RP0/CPU0:ios#show crypto key mypubkey rsa
Thu May  7 16:19:06.606 IST
Key label: tp
Type      : RSA General purpose
Size      : 2048
Created   : 16:18:49 IST Thu May 07 2020
Data      :
30820122 300D0609 2A864886 F70D0101 01050003 82010F00 3082010A 02820101
00CAC6E9 737D5ACF 31D0F8F2 281A450C 4F251D95 53587BCA 13592991 0AF2E6AF
02A89439 1DEDA683 C467C55F 032F05F3 A72DDED9 323F171A FDDEE3C4 DC124439
A78652F6 BB97BE63 F5AC8E3B 03B9B141 DD5D1AAE E41D7C15 28DE96E4 D3F4CE33
B12C477A 525CBDF6 17B92A8C E94A816E 7C4BCEFA 0EA7972D A3B0CBF1 A1DED71E
36CE08B7 3EF477A7 7B875BE1 E1B9E3A8 1E6C2717 6AB6D5AE 3A11D200 B32F4CCB
0B4163A4 E44D5729 70ECFEE6 4713D1CC 588C8AFB D3EE9891 B27BCA5B 8CD82B76
C278B32C 9A24B2EA 0CB9F2F3 6D1A1C95 044106F6 7E71520E B0201414 1E15B1C8
A88E4164 F3474B66 86CF4DFB 8B0DA66C 8C80C8BF EF192CC8 F85FBF71 D1A35B7A
05020301 0001

```

3. Declare a Certification Authority and configure a trustpoint.

configure

crypto ca trustpoint {*ca-name*}

enrollment url {*ca-url*}

subject-name {*x.500-name*}

serial-number

rsakeypair {*keypair-label*}

crl optional

ip-address *ip-address*

commit

```

RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios (config)# crypto ca trustpoint myca

RP/0/RP0/CPU0:ios (config-trustp)# enrollment url http://209.165.200.226

RP/0/RP0/CPU0:ios (config-trustp)# subject-name CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US

RP/0/RP0/CPU0:ios (config-trustp)#serial-number
RP/0/RP0/CPU0:ios (config-trustp)# rsa keypair tp

RP/0/RP0/CPU0:ios (config-trustp)# crl optional
RP/0/RP0/CPU0:ios (config-trustp)# ip-address 10.105.57.100

RP/0/RP0/CPU0:ios (config-trustp)# commit

```

4. Authenticate the CA—The router must authenticate the CA by obtaining the self-signed certificate of the CA, which contains the public key of the CA. Because the certificate of the CA is self-signed (the CA signs its own certificate), manually authenticate the public key of the CA by contacting the CA administrator to compare the fingerprint of the CA certificate.

crypto ca authenticate *ca-name*

```
RP/0/RP0/CPU0:ios#crypto ca authenticate myca
Thu May  7 16:20:08.458 IST
Serial Number  : 01
  CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US
Issued By      :
  CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US
Validity Start : 11:55:46 UTC Wed Jan 08 2020
Validity End   : 11:55:46 UTC Sat Jan 07 2023
SHA1 Fingerprint:
  70562AA850DE24B2D94AACF62528042E53C33D23
Do you accept this certificate? [yes/no]: yes
```

5. Request Device Certificates—You must obtain a signed certificate from the CA for each of your router's RSA key pairs.

crypto ca enroll ca-name

```
RP/0/RP0/CPU0:ios#crypto ca enroll myca
Thu May  7 16:20:34.776 IST
% Start certificate enrollment ...
% Create a challenge password. You will need to verbally provide this
  password to the CA Administrator in order to revoke your certificate.
% For security reasons your password will not be saved in the configuration.
% Please make a note of it.

Password:
Re-enter Password:

% The subject name in the certificate will include:
CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US
% The subject name in the certificate will include: myhost.mydomain.com
% The serial number in the certificate will be: 93f379c1
% The IP address in the certificate is 10.105.57.100
  Fingerprint: 41304434 42393333 45314143 42443134
```

6. Verify CA certificate.

show crypto ca certificates certificate-name

```
RP/0/RP0/CPU0:ios#show crypto ca certificates myca
Thu May  7 16:21:24.633 IST

Trustpoint      : myca
=====
CA certificate
Serial Number   : 01
Subject:
  CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US
Issued By       :
  CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US
Validity Start  : 11:55:46 UTC Wed Jan 08 2020
Validity End    : 11:55:46 UTC Sat Jan 07 2023
SHA1 Fingerprint:
  70562AA850DE24B2D94AACF62528042E53C33D23
Router certificate
Key usage       : General Purpose
Status          : Available
Serial Number   : 08:4D
Subject:
  serialNumber=93f379c1,unstructuredAddress=10.105.57.100,
unstructuredName=myhost.mydomain.com,CN=ncs,OU=BU,O=Govt,
L=Newyork,ST=NY,C=US
Issued By       :
  CN=ncs,OU=BU,O=Govt,L=Newyork,ST=NY,C=US
Validity Start  : 10:44:51 UTC Thu May 07 2020
```

```

Validity End      : 10:44:51 UTC Fri May 07 2021

CRL Distribution Point
    http://7200.cisco.com
SHA1 Fingerprint:
    C9454B6DD92A057A1DDB60740E0459243B070B24
Associated Trustpoint: myca

```

7. Configure IKEv2 profile.

config

ikev2 profile *profile-name*

match identity remote address *{ipv4-address [mask]}*

pki trustpoint *trustpoint-label*

lifetime *seconds*

authentication local rsa-signature

authentication remote rsa-signature

commit

```

RP/0/RP0/CPU0:ios#configure
Thu May  7 16:22:33.804 IST
RP/0/RP0/CPU0:ios(config)#ikev2 profile IP1
RP/0/RP0/CPU0:ios(config-ikev2-profile-IP1)#match identity remote address 10.1.1.2
255.255.255.255
RP/0/RP0/CPU0:ios(config-ikev2-profile-IP1)#pki trustpoint myca
RP/0/RP0/CPU0:ios(config-ikev2-profile-IP1)#lifetime 86400
Router(config)#
RP/0/RP0/CPU0:ios(config-ikev2-profile-IP1)#authentication local rsa-signature
RP/0/RP0/CPU0:ios(config-ikev2-profile-IP1)#authentication remote rsa-signature
RP/0/RP0/CPU0:ios(config-ikev2-profile-IP1)#commit

```

8. Configure the GCC2 interface. See [Configuring the GCC Interface, on page 105](#).

9. Configure OTNsec on ODU4 controller.

config

controller ODU4 *R/S/I/P*

otnsec

source ipv4 *ipv4-address*

destination ipv4 *ipv4-address*

session-id *session-id*

policy *policy-name*

ikev2 *profile-name*

```

RP/0/RP0/CPU0:ios#configure
Thu May  7 16:27:57.294 IST
RP/0/RP0/CPU0:ios(config)#controller ODU4 0/1/0/0/1
RP/0/RP0/CPU0:ios(config-odu4)#otnsec
RP/0/RP0/CPU0:ios(config-otnsec)#source ipv4 10.1.1.1
RP/0/RP0/CPU0:ios(config-otnsec)#destination ipv4 10.1.1.2
RP/0/RP0/CPU0:ios(config-otnsec)#session-id 1
RP/0/RP0/CPU0:ios(config-otnsec)#policy OP1

```

```
RP/0/RP0/CPU0:ios(config-otnsec)#ikev2 IP1
RP/0/RP0/CPU0:ios(config-otnsec)#commit
```

10. Verify the IKEv2 session.

```
RP/0/RP0/CPU0:ios#show ikev2 session
Wed Sep 22 21:09:38.363 IST
```

```
Session ID                               : 4
=====
Status                                   : UP-ACTIVE
IKE Count                                 : 1
Child Count                               : 1
IKE SA ID                                 : 373
-----
Local                                     : 10.1.1.1/500
Remote                                    : 10.1.1.2/500
Status(Description)                       : READY (Negotiation done)
Role                                       : Initiator

Child SA
-----
Local Selector                            : 10.1.1.1/1 - 10.1.1.1/1
Remote Selector                           : 10.1.1.2/1 - 10.1.1.2/1
ESP SPI IN/OUT                            : 0x2803 / 0x2800
```

```
RP/0/RP0/CPU0:ios#show ikev2 summary
Wed Sep 22 21:09:42.354 IST
```

```
IKEv2 SA Summary
-----
Total SA (Active/Negotiating)             : 1 (1/0)
Total Outgoing SA (Active/Negotiating)    : 1 (1/0)
Total Incoming SA (Active/Negotiating)    : 0 (0/0)
```

You May Be Interested In

- For more information about IKEv2, see [RFC 7296](#).
- For more information about NCS 1004, see the [NCS 1004 datasheet](#).



CHAPTER 8

GMPLS UNI for Packet and Optical Integration

With the cloud becoming increasingly central to business operations, packet and optical network services must evolve to become more efficient and dynamic. Closer integration of packet and optical networks becomes critical especially in the control plane.

- [Understanding GMPLS UNI, on page 117](#)
- [Use Case Overview, on page 118](#)
- [Prerequisites, on page 118](#)
- [Limitations, on page 119](#)
- [Configuration Workflow, on page 119](#)
- [Verification, on page 128](#)
- [General Troubleshooting, on page 135](#)
- [You May Be Also Interested In, on page 135](#)

Understanding GMPLS UNI

Generalized Multiprotocol Label Switching (GMPLS) User Network Interface (UNI) or GMPLS UNI is a key technology that enables this integration. GMPLS UNI enables packet networks to directly tap into the optical transport control plane to coordinate its resource requirements with the optical transport network. Leveraging open standards, GMPLS UNI optimizes network resources and improves network utilization across packet and optical networks.

Channel Spacing

DWDM grid in the optical spectrum can be divided into multiple channels so that each channel can carry traffic independently. The number of channels that we receive from the DWDM grid depends on the channel spacing. For example, the lower the channel spacing, the higher the number of channels, and also conversely.

GMPLS has two types of channel spacing:

- Fixed Grid channel spacing - The channel spacing is fixed to 50 GHz and supports 100 and 200-Gbps traffic.
- Flexible Grid channel spacing - The channel spacing is 6.25 GHz and supports all data rates.

The **neighbor flexi-grid-capable** command enables GMPLS UNI flexible grid channel spacing. This command is executed during the [LMP configuration](#) configuration.


Use Case Overview

GMPLS UNI technology addresses the following customer needs in packet and optical networks:

- Effective usage of the DWDM grid with minimal wastage of spectral bandwidth
- Transmission of mixed bit-rate or mixed modulation data in a grid with different channel widths

To address these needs, you create a tunnel between two NCS 1004 nodes to carry traffic using the GMPLS UNI technology as shown in the following figure.

Figure 6: GMPLS UNI Reference Model



UNI-C is the client or packet or router node; for example, NCS 1004 nodes. UNI-N is the network or optical node; for example, NCS 2000 nodes.

The Link Management Protocol (LMP) link is created to establish connectivity between a NCS 2000 node and a NCS 1004 node. The tunnel is then created between the trunk interfaces of the source and destination NCS 1004 nodes to carry traffic. When the tunnel is created between NCS 1004 nodes, a circuit is internally created between the NCS 2000 nodes. The circuit is created to perform path computation, restoration, and reversion functions.

The tunnel can be created between the source and destination NCS 1004 nodes without involving NCS 2000 nodes in the middle. However, the restoration and reversion capabilities are provided only by the NCS 2000 nodes using GMPLS UNI.

Prerequisites

Before you create a tunnel using GMPLS UNI, fulfill these prerequisites:

- NCS 1004 node must have both the MPLS and MPLS-TE packages. The package names are ncs1004-mpls and ncs1004-mpls-te-rsvp.
- NCS 2000 node must have a valid license for ROADM and WSON support.
- The management IP addresses of NCS 1004 and NCS 2000 nodes must be accessible.
- The administrative state of the trunk port of the optics controller on the NCS 1004 node must not be in the shutdown state.

Limitations

Configuration Workflow

Perform the following tasks in sequence to create a tunnel using GMPLS UNI:

Configurations on the NCS 2000 node:

1. GMPLS signaled LMP circuit creation.
 - [#unique_105](#)
2. [Retrieve Ifindex from NCS 2000 Node, on page 124](#)

Configurations on the node:

1. Configure LMP on Cisco NCS 1002 Node.
 - [Configure LMP on Cisco NCS 1004 Node , on page 125](#)
2. [Configure RSVP on NCS 1004 Node, on page 126](#)
3. Configure MPLS Tunnel on a NCS 1002 Node.
 - [Configure MPLS Tunnel on a NCS 1004 Node for Numbered Circuit, on page 127](#)

Configure an LMP link and Alien Wavelength on an NCS 2000 node using CTC

Establish connectivity and enable advanced opUse this task to set up a static LMP link between a Cisco NCS 2000 node and either an NCS 1004 node. The CTC LMP creation wizard enables you to select source and destination endpoints, configure optical parameters, and set alien wavelength options. tical parameters and alien wavelength configurations for signaled numbered circuits between Cisco optical network nodes.

Use this task to set up a static LMP link between a Cisco NCS 2000 node and either an NCS 1002 or NCS 1004 node. The CTC LMP creation wizard enables you to select source and destination endpoints, configure optical parameters, and set alien wavelength options.

Procedure

- Step 1** From the **View** menu, choose **Go to Network View**.
- Step 2** Click the **Provisioning > LMP** tabs.
- Step 3** Click **Create**.
The LMP Creation window appears.
- Step 4** Click **Signaled** in the **Router Not Managed by CTC** area.
A wizard appears with the following options:

LMP Origination, LMP Termination, Optical Parameters, and Alien Wavelength**Step 5**

In the LMP Origination screen of the wizard, provision these parameters:

- From the **Originating Node** drop-down list, choose the source node of the LMP.
If the source node is Cisco NCS 1004, the destination node must be MSTP, and the other way round.
- From the **Local Interfaces** drop-down list, choose an available interface.
- Choose the Type, Shelf, Slot, and Port for Ingress Port Selection and Egress Port Selection.
- Choose **Numbered** interface.
- Enter the IP address of the source node in the **Interface IP** field.
- Set the mode of revertive restoration to either UNI-C or UNI-N. If the mode is set to UNI-C, the reversion of the circuit from the restored path to the original path is initiated by the UNI client that is connected to NCS 1004. If the mode is set to UNI-N, the reversion of the circuit is initiated by the DWDM network and can either be a manual revert or an auto revert.
- Enter the RSVP signaling interval and RSVP signaling missed values in the respective fields.
- Click **Next**.

Step 6

In the LMP Termination screen of the wizard, provision these parameters:

- From the **Terminating Node** drop-down list, choose the destination node of the LMP; for example, MSTP node.
- From the **Rx Port Selection** area, perform the following.
 - Choose the card type from the **Type** drop-down list.
 - Choose a shelf from the **Shelf** drop-down list.
 - Choose a source slot from the **Slot** drop-down list
 - Choose a port from the **Port** drop-down list.
- From the **Tx Port Selection** area, perform the following.
 - Choose the card type from the **Type** drop-down list.
 - Choose a shelf from the **Shelf** drop-down list.
 - Choose a destination slot from the **Slot** drop-down list.
 - Choose a port from the **Port** drop-down list
- Enter the IP address of the destination node in the **Interface IP** field.
- Set the mode of revertive restoration to either UNI-C or UNI-N. If the mode is set to UNI-C, the reversion of the circuit from the restored path to the original path is initiated by the UNI client that is connected. If the mode is set to UNI-N, the reversion of the circuit is initiated by the DWDM network and can be either a manual revert or an auto revert.
- Enter the remote Ifindex of NCS 1004 node (in decimals) in the **Remote If Index** field.
- Click **Next**.

- Step 7** In the Optical Parameters screen of the wizard, provision these parameters:
- **Allow Regeneration**—When checked, the computed path traverses through the regeneration site only if the optical validation is not satisfied. You can regenerate a circuit that is created from the UNI interface. If a transparent path is feasible, the regenerator is not used.
 - **UNI State**—Choose **Enable** or **Disable** from the UNI State drop-down list.

The Enable state is used to configure the UNI interface for the circuits to pass through, between the router and the DWDM node. In the Disable state, the interface is configured but not active, and so the circuit activation is rejected. When the status is changed from Enable to Disable, all active circuits on the interface are deleted.
 - **Description**—Enter the description of the UNI interface. The description can be up to 256 characters.
 - **Label**—Enter an alphanumeric string. This label is a unique circuit identifier.
 - **Validation**—Sets the optical validation mode.
 - **Full**—The circuit is created when the circuit validation result is greater than or equal to the acceptance threshold value.
 - **None**—The circuit is created without considering the acceptance threshold value. The Opt Valid column in the Circuits tab displays the value as **Not Valid**.
 - **Inherited**—The restoration circuit inherits the validation and acceptance threshold values from the primary circuit.
 - **Acceptance threshold**—Sets the acceptance threshold value for the GMPLS circuit. The circuit is created if the actual acceptance threshold value is greater than, or equal to, the value set in this field.
 - Green—Indicates that the channel failure risk is 0%.
 - Yellow—Indicates that the channel failure risk is between 0% and 16%.
 - Orange—Indicates that the channel failure risk is between 16% and 50%.
 - Red—Indicates that the channel failure risk is greater than 50%.
 - **Restoration**—Check this check box to enable the restoration of the GMPLS circuits on the UNI interface.
 - **Revert**—Check this check box to enable the revert of the GMPLS circuits on the UNI interface.
 - **Auto Revert**—Click this radio button to automatically revert the circuit from the restored path to the original path after the failure is fixed, WSON alarms are acknowledged, and the soak time expires.
 - **Manual Revert**—Click this radio button to manually revert the circuit from the restored path to the original path after the failure is fixed, the WSON alarms are acknowledged, and the soak time expires.
 - **Soak Time**—Enter the time (in hours, minutes, and seconds) in the Soak Time field that the circuit on the restored path waits before moving to the original path after the failure is fixed. The circuit reverts to the original path after the soak time expires. The soak time must be set only if both the **Restoration** and **Revert** check boxes are checked.
- Step 8** Click **Next**.
- Step 9** In the Alien wavelength screen of the wizard, provision these parameters.
- From the **Alien Wavelength** drop-down list, choose the alien wavelength class.

- From the **Trunk Selection** drop-down list, choose 100G, 200G, or 250G.
- From the **FEC** drop-down list, choose a valid value for forward error correction (FEC) mode. If an invalid FEC value is chosen, LMP link is created; however, the circuit creation fails.
- Click **Finish** to create an LMP link.

The LMP link with specified optical parameters and alien wavelength settings is created and listed in CTC, ready for signaled numbered circuit provisioning.

Configure an LMP link with Alien Wavelength on NCS 2000 using CTC for signaled unnumbered circuits

Establish connectivity between NCS 2000 and NCS 1002/1004 nodes using a signaled, unnumbered LMP link with Alien Wavelength support.

Use the LMP creation wizard in Cisco Transport Controller (CTC) to provision link parameters, optical settings, and Alien Wavelength options for interoperability between managed and unmanaged router nodes.

Procedure

- Step 1** From the **View** menu, choose **Go to Network View**.
- Step 2** Click the **Provisioning > LMP** tabs.
- Step 3** Click **Create**.
- The LMP Creation window appears.
- Step 4** Click **Signaled** in the **Router Not Managed by CTC** area.
- A wizard appears with options for **LMP Origination**, **LMP Termination**, **Optical Parameters**, and **Alien Wavelength**.
- Step 5** On the **LMP Origination** screen, provision these parameters:
- From the **Originating Node** drop-down list, choose the source node of the LMP.
 - From the **Local Interfaces** drop-down list, choose an available interface.
 - Choose the Type, Unit, and Port for Ingress Port Selection and Egress Port Selection.
 - Choose **Unnumbered** interface.
 - The IP address of the source node selected appears in the **IP** field.
 - Set the mode of revertive restoration to UNI-N. If the mode is set to UNI-N, the reversion of the circuit is initiated by the DWDM network and can either be a manual revert or an auto revert.
 - Click **Next**.
- Step 6** On the LMP Termination screen of the wizard, provision these parameters:

- From **Interfaces Configuration**:
 - Enter the NCS 1004 system IP address in the **System IP** field.
- Enter the IP address of the source node in the **Communication Channel** field.
- Enter the SNMP Ifindex value of optic trunk in the **Remote If Index** field.
- Click **Next**.

Step 7 In the Optical Parameters screen of the wizard, provision these parameters:

- **Allow Regeneration**—When checked, the computed path traverses through the regeneration site only if the optical validation is not satisfied. You can regenerate a circuit that is created from the UNI interface. If a transparent path is feasible, the regenerator is not used.
- **UNI State**—Choose **Enable** or **Disable** from the UNI State drop-down list.

The Enable state is used to configure the UNI interface for the circuits to pass through, between the router and the DWDM node. In the Disable state, the interface is configured but not active, and so the circuit activation is rejected. When the status is changed from Enable to Disable, all active circuits on the interface are deleted.
- **Description**—Enter the description of the UNI interface like **Signal Unnumb LMP**. The description can be up to 256 characters.
- **Label**—Enter an alphanumeric string. This label is a unique circuit identifier.
- **Validation**—Sets the optical validation mode.
 - **Full**—The circuit is created when the circuit validation result is greater than or equal to the acceptance threshold value.
 - **None**—The circuit is created without considering the acceptance threshold value. The Opt Valid column in the Circuits tab displays the value as **Not Valid**.
 - **Inherited**—The restoration circuit inherits the validation and acceptance threshold values from the primary circuit.
- **Acceptance Threshold**—Sets the acceptance threshold value for the GMPLS circuit. The circuit is created if the actual acceptance threshold value is greater than, or equal to, the value set in this field.
 - **Green**—Indicates that the channel failure risk is 0%.
 - **Yellow**—Indicates that the channel failure risk is between 0% and 16%.
 - **Orange**—Indicates that the channel failure risk is between 16% and 50%.
 - **Red**—Indicates that the channel failure risk is greater than 50%.
- **Restoration**—Check this check box to enable the restoration of the GMPLS circuits on the UNI interface.
- **Revert**—Check this check box to enable the revert of the GMPLS circuits on the UNI interface.
- **Auto Revert**—Click this radio button to automatically revert the circuit from the restored path to the original path after the failure is fixed, WSON alarms are acknowledged, and the soak time expires.
- **Manual Revert**—Click this radio button to manually revert the circuit from the restored path to the original path after the failure is fixed, the WSON alarms are acknowledged, and the soak time expires.

- **Soak Time**—Enter the time (in hours, minutes, and seconds) in the Soak Time field that the circuit on the restored path waits before moving to the original path after the failure is fixed. The circuit reverts to the original path after the soak time expires. The soak time must be set only if both the **Restoration** and **Revert** check boxes are checked.

Step 8 Click **Next**.

Step 9 In the Alien wavelength screen of the wizard, provision these parameters.

- From the **Alien Wavelength** drop-down list, choose the alien wavelength class such as NCS 1004.
- From the **Trunk Selection** drop-down list, choose 100G, 200G, or 250G.
- From the **FEC** drop-down list, choose a valid value for forward error correction (FEC) mode. If an invalid FEC value is chosen, LMP link is created; however, the circuit creation fails.
- Click **Finish** to create an LMP link.

The newly created signaled LMP unnumbered circuit link appears in the LMP table in CTC.

The LMP link between NCS 2000 and NCS 1002/1004 is established, supporting signaled unnumbered circuits with full optical and Alien Wavelength configuration.

Retrieve Ifindex from NCS 2000 Node

The Ifindex value of all the LMP ports of NCS 2000 node can be retrieved using CTC or TL1.

Using CTC:

From the **Provisioning > LMP** tab, retrieve the Ifindex value in decimal format under the **Originating Interface Index** column.

This Ifindex value is used in the **neighbor interface-id unnumbered** command during the [LMP configuration](#) configuration.

Using TL1:

1. Log in to the TL1 interface and issue the following command.
2. **rtrv-unicfg ::all:1;**

This command retrieves the Ifindex of all the LMP ports of NCS 2000 node in hexadecimal format. This must be converted to decimal format and used in remote Ifindex of NCS 1004 node during the [LMP configuration](#).

TL1 Output

```
PSLINE-81-1-9-RX:PSLINE-81-1-9-TX,10.77.142.92,3.3.3.4,3.3.3.3,0.0.0.0,VALMODE=NONE,ADMINSTATE=UP,
RESTTYPE=REVERT,USPWROFS=0.0,
DSPWROFS=0.0,ALLOWREGEN=NO,UNICTRLMODE=CLIENT,REVERTMODE=MANUAL,SOAK=00-01-00,
RESTVALMODE=NONE,TERMINTFDX=0,ORIGINTFIDX=7f000d12,NUMBERED=TRUE,UNIMODE=GMPLS
```

```
PSLINE-81-1-10-RX:PSLINE-81-1-10-TX,10.77.142.92,4.4.4.4,4.4.4.3,0.0.0.0,VALMODE=NONE,ADMINSTATE=UP,
RESTTYPE=REVERT,USPWROFS=0.0,DSPWROFS=0.0,ALLOWREGEN=NO,UNICTRLMODE=CLIENT,
REVERTMODE=MANUAL,SOAK=00-01-00,RESTVALMODE=NONE,TERMINTFDX=0,
ORIGINTFIDX=7f000d14,NUMBERED=TRUE,UNIMODE=GMPLS
```

The Ifindex of port 81-1-9 is 7f000d12 (in hexadecimal) and 2130709778 (in decimal). The Ifindex of port 81-1-10 is 7f000d14 (in hexadecimal) and 2130709780 (in decimal).

Configure LMP on Cisco NCS 1004 Node

LMP is a logical link that is created on the trunk optics controller of the source and destination NCS 1004 nodes of the tunnel.

configure

lmp

gmpls optical-uni

controller optics *Rack/Slot/Instance/Port*

neighbor *name*

neighbor link-id ipv4 unicast *ipv4-address*

neighbor flexi-grid-capable

neighbor interface-id unnumbered *interface-id*

link-id ipv4 unicast *ipv4-address*

router-id ipv4 unicast *ipv4-address*

commit

Important Notes

- **neighbor link-id ipv4 unicast** *ipv4-address* is the IP address of the MSTP interface on the NCS 2000 node.
- **neighbor flexi-grid-capable** enables GMPLS UNI flexible grid channel spacing.
- **neighbor interface-id unnumbered** *interface-id* is the optical interface ID of the neighbor. This value is the Ifindex value of all the LMP ports of NCS 2000 node in decimal format that is manually retrieved from CTC or TL1. See [Retrieve Ifindex from NCS 2000 Node, on page 124](#) to retrieve the Ifindex.
- **link-id ipv4 unicast** *ipv4-address* is the IP address of the optics controller on the current NCS 1004 node.
- **router-id ipv4 unicast** *ipv4-address* is the neighbor router IP address for GMPLS UNI.

Running Configuration

The following is a sample of configuring LMP on the source NCS 1004 node.

show running-config lmp

```
Mon Jul  1 14:42:46.856 IST
lmp
gmpls optical-uni
  controller Optics0/0/0/0
  neighbor ncs1k
  neighbor link-id ipv4 unicast 10.1.1.1
  neighbor flexi-grid-capable
```

```

neighbor interface-id unnumbered 2130706976
link-id ipv4 unicast 10.0.1.1
!
controller Optics0/0/0/1
neighbor ncs1k
neighbor link-id ipv4 unicast 10.1.3.3
neighbor flexi-grid-capable
neighbor interface-id unnumbered 2130707232
link-id ipv4 unicast 10.0.3.3
!
controller Optics0/1/0/0
neighbor ncs1k
neighbor link-id ipv4 unicast 10.1.4.4
neighbor flexi-grid-capable
neighbor interface-id unnumbered 2130706964
link-id ipv4 unicast 10.0.4.4
!
controller Optics0/1/0/1
neighbor ncs1k
neighbor link-id ipv4 unicast 10.1.5.5
neighbor flexi-grid-capable
neighbor interface-id unnumbered 2130706966
link-id ipv4 unicast 10.0.5.5
!
neighbor ncs1k
ipcc routed
router-id ipv4 unicast 10.127.60.48
!
router-id ipv4 unicast 10.105.57.101
!
!
!

```

The following sample shows the brief summary of the tunnel status and configuration.

show mpls traffic-eng tunnels optical-uni brief

Wed Sep 22 17:08:13.132 IST

TUNNEL NAME	DESTINATION	STATUS	STATE
GMPLS-UNI-Optics0/3/0/1	10.24.1.1	up	up
GMPLS-UNI-Optics0/0/0/1	10.34.1.1	up	up

Displayed 2 (of 2) heads, 0 (of 0) midpoints, 0 (of 0) tails
 Displayed 2 up, 0 down, 0 recovering, 0 recovered heads

Configure RSVP on NCS 1004 Node

Resource Reservation Protocol (RSVP) with an appropriate timeout must be configured on the source and destination NCS 1004 nodes of the tunnel.

configure

rsvp

controller optics *Rack/Slot/Instance/Port*

signalling refresh out-of-band interval *interval*

signalling refresh out-of-band missed *mis-count*

commit

The following is a sample of configuring RSVP on the source NCS 1004 node.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#rsvp
RP/0/RP0/CPU0:ios(config-rsvp)#controller optics 0/0/0/6
RP/0/RP0/CPU0:ios(config-rsvp-cntl)#signalling refresh out-of-band interval 3600
RP/0/RP0/CPU0:ios(config-rsvp-cntl)#signalling refresh out-of-band missed 24
RP/0/RP0/CPU0:ios(config-rsvp-cntl)#commit
```

Configure MPLS Tunnel on a NCS 1004 Node for Numbered Circuit

Ensure that the administrative state of the trunk port of the optics controller on the NCS 1004 node is not in shutdown state.

```
configure
mpls traffic-eng
gmpls optical-uni
controller optics Rack/Slot/Instance/Port
tunnel-properties
tunnel-id id
destination ipv4 unicast ipv4-address
path-option 10 no-ero lockdown
commit
```

Important Notes

- **destination ipv4 unicast** *ipv4-address* is the IP address of the optics controller on the destination NCS 1004 node.
- Explicit Route Object (ERO) - Includes one or more routes to use from a list of specified nodes for a tunnel.
- Exclude Route Object (XRO) - Excludes one or more routes to use from a list of specified nodes for a tunnel.

Running Configuration

The following is a sample of configuring the MPLS tunnel on the source NCS 1004 node.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#mpls traffic-eng
RP/0/RP0/CPU0:ios(config-mpls-te)#gmpls optical-uni
RP/0/RP0/CPU0:ios(config-te-gmpls-uni)#controller optics 0/0/0/6
RP/0/RP0/CPU0:ios(config-te-gmpls-cntl)#tunnel-properties
RP/0/RP0/CPU0:ios(config-te-gmpls-tun)#tunnel-id 100
RP/0/RP0/CPU0:ios(config-te-gmpls-tun)#destination ipv4 unicast 10.20.20.20
RP/0/RP0/CPU0:ios(config-te-gmpls-tun)#path-option 10 no-ero lockdown
RP/0/RP0/CPU0:ios(config-te-gmpls-tun)#commit
```

The following is a sample of configuring the MPLS tunnel on the destination NCS 1004 node.

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

```

RP/0/RP0/CPU0:ios(config)#mpls traffic-eng
RP/0/RP0/CPU0:ios(config-mpls-te)#gmpls optical-uni
RP/0/RP0/CPU0:ios(config-te-gmpls-uni)#controller optics 0/0/0/6
RP/0/RP0/CPU0:ios(config-te-gmpls-uni)#commit

```

Verification

Use the show commands in the following table to verify the GMPLS UNI tunnel, RSVP, and LMP configuration.

Table 11: Show Commands

Show Commands	Description
show mpls traffic-eng link-management optical-uni controller optics	Displays detailed GMPLS information of a specific optics controller.
show mpls traffic-eng link-management optical-uni	Displays detailed GMPLS information of all the optics controllers.
show mpls traffic-eng tunnels	Displays information about tunnels.
show mpls traffic-eng link-management optical-uni tabular	Displays detailed GMPLS information of all the optics controllers in tabular format.
show mpls traffic-eng tunnels tabular	Displays information about all the tunnels in tabular format.
show lmp gmpls optical-uni	Verifies LMP configuration and state.
show rsvp neighbors	Displays information about RSVP neighbors.

Sample Outputs

show mpls traffic-eng link-management optical-uni controller optics 0/0/0/13

Displays detailed GMPLS information of a specific optics controller.

```

Mon Jul 1 20:05:27.209 IST
Optical interface: Optics0/0/0/0
Overview:
  IM state: Up
  Child interface: : IM state Unknown
  OLM/LMP state: Up
  Optical tunnel state: up
Connection:
  Tunnel role: Tail
  Tunnel-id: 15, LSP-id 3, Extended tunnel-id 10.105.57.100
  Tunnel source: 10.105.57.100, destination: 10.11.1.1
  Optical router-ids: Local: 10.105.57.101, Remote: 10.127.60.48
  Label source: UNI-N
Upstream label:
  Optical label:
  Grid           : DWDM
  Channel spacing : 6.25 GHz
  Identifier      : 0
  Channel Number  : -277

```

```

Downstream label:
  Optical label:
    Grid           : DWDM
    Channel spacing : 6.25 GHz
    Identifier      : 0
    Channel Number  : -277
SRLG discovery: Disabled
SRLG announcement: None
Switching Type: lsc
MTU: 9212
Admission Control:
  Upstream: Admitted (LSP ID: 3)
  Downstream: Admitted (LSP ID: 3)
OLM/LMP adjacency information:
  Adjacency status: Up
  Local:
    node ID: 10.105.57.101
    link interface ID: 10
    link ID: 10.11.1.1
  Neighbor:
    node ID: 10.127.60.48 (VEGA2K-Site-3_48)
    link interface ID: 2130706976
    link ID: 10.1.1.1
    IPCC: Routed to 10.127.60.48
Optical capabilities:
  Controller type: DWDM
  Channel spacing: 6.25 GHz
  Default channel: 0
  784 supported channels:
    -303, -302, -301, -300, -299, -298, -297, -296
    -295, -294, -293, -292, -291, -290, -289, -288
    -287, -286, -285, -284, -283, -282, -281, -280
    -279, -278, -277, -276, -275, -274, -273, -272
    -271, -270, -269, -268, -267, -266, -265, -264
    -263, -262, -261, -260, -259, -258, -257, -256
    -255, -254, -253, -252, -251, -250, -249, -248
    -247, -246, -245, -244, -243, -242, -241, -240
    -239, -238, -237, -236, -235, -234, -233, -232
    -231, -230, -229, -228, -227, -226, -225, -224
    -223, -222, -221, -220, -219, -218, -217, -216
    -215, -214, -213, -212, -211, -210, -209, -208
    -207, -206, -205, -204, -203, -202, -201, -200
    -199, -198, -197, -196, -195, -194, -193, -192
    -191, -190, -189, -188, -187, -186, -185, -184
    -183, -182, -181, -180, -179, -178, -177, -176
    -175, -174, -173, -172, -171, -170, -169, -168
    -167, -166, -165, -164, -163, -162, -161, -160
    -159, -158, -157, -156, -155, -154, -153, -152
    -151, -150, -149, -148, -147, -146, -145, -144
    -143, -142, -141, -140, -139, -138, -137, -136
    -135, -134, -133, -132, -131, -130, -129, -128
    -127, -126, -125, -124, -123, -122, -121, -120
    -119, -118, -117, -116, -115, -114, -113, -112
    -111, -110, -109, -108, -107, -106, -105, -104
    -103, -102, -101, -100, -99, -98, -97, -96
    -95, -94, -93, -92, -91, -90, -89, -88
    -87, -86, -85, -84, -83, -82, -81, -80
    -79, -78, -77, -76, -75, -74, -73, -72
    -71, -70, -69, -68, -67, -66, -65, -64
    -63, -62, -61, -60, -59, -58, -57, -56
    -55, -54, -53, -52, -51, -50, -49, -48
    -47, -46, -45, -44, -43, -42, -41, -40
    -39, -38, -37, -36, -35, -34, -33, -32
    -31, -30, -29, -28, -27, -26, -25, -24

```

-23, -22, -21, -20, -19, -18, -17, -16
-15, -14, -13, -12, -11, -10, -9, -8
-7, -6, -5, -4, -3, -2, -1, 0
1, 2, 3, 4, 5, 6, 7, 8
9, 10, 11, 12, 13, 14, 15, 16
17, 18, 19, 20, 21, 22, 23, 24
25, 26, 27, 28, 29, 30, 31, 32
33, 34, 35, 36, 37, 38, 39, 40
41, 42, 43, 44, 45, 46, 47, 48
49, 50, 51, 52, 53, 54, 55, 56
57, 58, 59, 60, 61, 62, 63, 64
65, 66, 67, 68, 69, 70, 71, 72
73, 74, 75, 76, 77, 78, 79, 80
81, 82, 83, 84, 85, 86, 87, 88
89, 90, 91, 92, 93, 94, 95, 96
97, 98, 99, 100, 101, 102, 103, 104
105, 106, 107, 108, 109, 110, 111, 112
113, 114, 115, 116, 117, 118, 119, 120
121, 122, 123, 124, 125, 126, 127, 128
129, 130, 131, 132, 133, 134, 135, 136
137, 138, 139, 140, 141, 142, 143, 144
145, 146, 147, 148, 149, 150, 151, 152
153, 154, 155, 156, 157, 158, 159, 160
161, 162, 163, 164, 165, 166, 167, 168
169, 170, 171, 172, 173, 174, 175, 176
177, 178, 179, 180, 181, 182, 183, 184
185, 186, 187, 188, 189, 190, 191, 192
193, 194, 195, 196, 197, 198, 199, 200
201, 202, 203, 204, 205, 206, 207, 208
209, 210, 211, 212, 213, 214, 215, 216
217, 218, 219, 220, 221, 222, 223, 224
225, 226, 227, 228, 229, 230, 231, 232
233, 234, 235, 236, 237, 238, 239, 240
241, 242, 243, 244, 245, 246, 247, 248
249, 250, 251, 252, 253, 254, 255, 256
257, 258, 259, 260, 261, 262, 263, 264
265, 266, 267, 268, 269, 270, 271, 272
273, 274, 275, 276, 277, 278, 279, 280
281, 282, 283, 284, 285, 286, 287, 288
289, 290, 291, 292, 293, 294, 295, 296
297, 298, 299, 300, 301, 302, 303, 304
305, 306, 307, 308, 309, 310, 311, 312
313, 314, 315, 316, 317, 318, 319, 320
321, 322, 323, 324, 325, 326, 327, 328
329, 330, 331, 332, 333, 334, 335, 336
337, 338, 339, 340, 341, 342, 343, 344
345, 346, 347, 348, 349, 350, 351, 352
353, 354, 355, 356, 357, 358, 359, 360
361, 362, 363, 364, 365, 366, 367, 368
369, 370, 371, 372, 373, 374, 375, 376
377, 378, 379, 380, 381, 382, 383, 384
385, 386, 387, 388, 389, 390, 391, 392
393, 394, 395, 396, 397, 398, 399, 400
401, 402, 403, 404, 405, 406, 407, 408
409, 410, 411, 412, 413, 414, 415, 416
417, 418, 419, 420, 421, 422, 423, 424
425, 426, 427, 428, 429, 430, 431, 432
433, 434, 435, 436, 437, 438, 439, 440
441, 442, 443, 444, 445, 446, 447, 448
449, 450, 451, 452, 453, 454, 455, 456
457, 458, 459, 460, 461, 462, 463, 464
465, 466, 467, 468, 469, 470, 471, 472
473, 474, 475, 476, 477, 478, 479, 480

```

Controller SRLGs
  None

```

show mpls traffic-eng link-management optical-uni

Displays detailed GMPLS information of all the optics controllers. MPLS tunnels are not created when the optics controller is in the shutdown state. The state is displayed as **Admin down**. Enter the **no shutdown** command under the optics controller to initiate the tunnel creation.

```
Mon Jul 1 20:00:42.108 IST
```

System Information:

```
Optical Links Count: 1 (Maximum Links Supported 100)
```

```
Optical interface: Optics0/0/0/0
```

Overview:

```

IM state: Up
Child interface: : IM state Unknown
OLM/LMP state: Up
Optical tunnel state: up

```

Connection:

```

Tunnel role: Tail
Tunnel-id: 15, LSP-id 3, Extended tunnel-id 10.105.57.100
Tunnel source: 10.105.57.100, destination: 10.11.1.1
Optical router-ids: Local: 10.105.57.101, Remote: 10.127.60.48
Label source: UNI-N

```

Upstream label:

```

Optical label:
  Grid           : DWDM
  Channel spacing : 6.25 GHz
  Identifier      : 0
  Channel Number  : -277

```

Downstream label:

```

Optical label:
  Grid           : DWDM
  Channel spacing : 6.25 GHz
  Identifier      : 0
  Channel Number  : -277

```

```
SRLG discovery: Disabled
```

```
SRLG announcement: None
```

```
Switching Type: lsc
```

```
MTU: 9212
```

Admission Control:

```

Upstream: Admitted (LSP ID: 3)
Downstream: Admitted (LSP ID: 3)

```

OLM/LMP adjacency information:

```
Adjacency status: Up
```

Local:

```

node ID: 10.105.57.101
link interface ID: 10
link ID: 10.11.1.1

```

Neighbor:

```

node ID: 10.127.60.48 (VEGA2K-Site-3_48)
link interface ID: 2130706976
link ID: 10.1.1.1
IPCC: Routed to 10.127.60.48

```

Optical capabilities:

```
Controller type: DWDM
```

```
Channel spacing: 6.25 GHz
```

```
Default channel: 0
```

```
784 supported channels:
```

```

-303, -302, -301, -300, -299, -298, -297, -296
-295, -294, -293, -292, -291, -290, -289, -288

```

-287, -286, -285, -284, -283, -282, -281, -280
 -279, -278, -277, -276, -275, -274, -273, -272
 -271, -270, -269, -268, -267, -266, -265, -264
 -263, -262, -261, -260, -259, -258, -257, -256
 -255, -254, -253, -252, -251, -250, -249, -248
 -247, -246, -245, -244, -243, -242, -241, -240
 -239, -238, -237, -236, -235, -234, -233, -232
 -231, -230, -229, -228, -227, -226, -225, -224
 -223, -222, -221, -220, -219, -218, -217, -216
 -215, -214, -213, -212, -211, -210, -209, -208
 -207, -206, -205, -204, -203, -202, -201, -200
 -199, -198, -197, -196, -195, -194, -193, -192
 -191, -190, -189, -188, -187, -186, -185, -184
 -183, -182, -181, -180, -179, -178, -177, -176
 -175, -174, -173, -172, -171, -170, -169, -168
 -167, -166, -165, -164, -163, -162, -161, -160
 -159, -158, -157, -156, -155, -154, -153, -152
 -151, -150, -149, -148, -147, -146, -145, -144
 -143, -142, -141, -140, -139, -138, -137, -136
 -135, -134, -133, -132, -131, -130, -129, -128
 -127, -126, -125, -124, -123, -122, -121, -120
 -119, -118, -117, -116, -115, -114, -113, -112
 -111, -110, -109, -108, -107, -106, -105, -104
 -103, -102, -101, -100, -99, -98, -97, -96
 -95, -94, -93, -92, -91, -90, -89, -88
 -87, -86, -85, -84, -83, -82, -81, -80
 -79, -78, -77, -76, -75, -74, -73, -72
 -71, -70, -69, -68, -67, -66, -65, -64
 -63, -62, -61, -60, -59, -58, -57, -56
 -55, -54, -53, -52, -51, -50, -49, -48
 -47, -46, -45, -44, -43, -42, -41, -40
 -39, -38, -37, -36, -35, -34, -33, -32
 -31, -30, -29, -28, -27, -26, -25, -24
 -23, -22, -21, -20, -19, -18, -17, -16
 -15, -14, -13, -12, -11, -10, -9, -8
 -7, -6, -5, -4, -3, -2, -1, 0
 1, 2, 3, 4, 5, 6, 7, 8
 9, 10, 11, 12, 13, 14, 15, 16
 17, 18, 19, 20, 21, 22, 23, 24
 25, 26, 27, 28, 29, 30, 31, 32
 33, 34, 35, 36, 37, 38, 39, 40
 41, 42, 43, 44, 45, 46, 47, 48
 49, 50, 51, 52, 53, 54, 55, 56
 57, 58, 59, 60, 61, 62, 63, 64
 65, 66, 67, 68, 69, 70, 71, 72
 73, 74, 75, 76, 77, 78, 79, 80
 81, 82, 83, 84, 85, 86, 87, 88
 89, 90, 91, 92, 93, 94, 95, 96
 97, 98, 99, 100, 101, 102, 103, 104
 105, 106, 107, 108, 109, 110, 111, 112
 113, 114, 115, 116, 117, 118, 119, 120
 121, 122, 123, 124, 125, 126, 127, 128
 129, 130, 131, 132, 133, 134, 135, 136
 137, 138, 139, 140, 141, 142, 143, 144
 145, 146, 147, 148, 149, 150, 151, 152
 153, 154, 155, 156, 157, 158, 159, 160
 161, 162, 163, 164, 165, 166, 167, 168
 169, 170, 171, 172, 173, 174, 175, 176
 177, 178, 179, 180, 181, 182, 183, 184
 185, 186, 187, 188, 189, 190, 191, 192
 193, 194, 195, 196, 197, 198, 199, 200
 201, 202, 203, 204, 205, 206, 207, 208
 209, 210, 211, 212, 213, 214, 215, 216
 217, 218, 219, 220, 221, 222, 223, 224

```

225, 226, 227, 228, 229, 230, 231, 232
233, 234, 235, 236, 237, 238, 239, 240
241, 242, 243, 244, 245, 246, 247, 248
249, 250, 251, 252, 253, 254, 255, 256
257, 258, 259, 260, 261, 262, 263, 264
265, 266, 267, 268, 269, 270, 271, 272
273, 274, 275, 276, 277, 278, 279, 280
281, 282, 283, 284, 285, 286, 287, 288
289, 290, 291, 292, 293, 294, 295, 296
297, 298, 299, 300, 301, 302, 303, 304
305, 306, 307, 308, 309, 310, 311, 312
313, 314, 315, 316, 317, 318, 319, 320
321, 322, 323, 324, 325, 326, 327, 328
329, 330, 331, 332, 333, 334, 335, 336
337, 338, 339, 340, 341, 342, 343, 344
345, 346, 347, 348, 349, 350, 351, 352
353, 354, 355, 356, 357, 358, 359, 360
361, 362, 363, 364, 365, 366, 367, 368
369, 370, 371, 372, 373, 374, 375, 376
377, 378, 379, 380, 381, 382, 383, 384
385, 386, 387, 388, 389, 390, 391, 392
393, 394, 395, 396, 397, 398, 399, 400
401, 402, 403, 404, 405, 406, 407, 408
409, 410, 411, 412, 413, 414, 415, 416
417, 418, 419, 420, 421, 422, 423, 424
425, 426, 427, 428, 429, 430, 431, 432
433, 434, 435, 436, 437, 438, 439, 440
441, 442, 443, 444, 445, 446, 447, 448
449, 450, 451, 452, 453, 454, 455, 456
457, 458, 459, 460, 461, 462, 463, 464
465, 466, 467, 468, 469, 470, 471, 472
473, 474, 475, 476, 477, 478, 479, 480
Controller SRLGs
None
    
```

show mpls traffic-eng link-management optical-uni tabular

Displays detailed GMPLS information of all the optics controllers in tabular format.

Mon Jul 1 15:10:50.472 IST

System Information:

Optical Links Count: 4 (Maximum Links Supported 100)

Interface	State		LMP adjacency	GMPLS tunnel		
	Admin	Oper		role	tun-id	state
Op0/0/0/0	up	up	up	Tail	15	up
Op0/0/0/1	up	up	up	Tail	16	up
Op0/1/0/0	up	up	up	Tail	17	up
Op0/1/0/1	up	up	up	Tail	18	up

show mpls traffic-eng tunnels

Displays information about tunnels.

Mon Jul 1 15:03:58.490 IST

LSP Tunnel 10.105.57.100 15 [5] is signalled, Signaling State: up

Tunnel Name: ckt0/0/0/0 Tunnel Role: Tail

Upstream label:

Optical label:

Grid : DWDM

```

Channel spacing      : 6.25 GHz
Identifier          : 0
Channel Number      : -277
Downstream label:
Optical label:
Grid               : DWDM
Channel spacing     : 6.25 GHz
Identifier          : 0
Channel Number      : -277
Signalling Info:
Src 10.105.57.100 Dst 10.11.1.1, Tun ID 15, Tun Inst 5, Ext ID 10.105.57.100
Router-IDs: upstream 10.127.60.48
             local   10.105.57.101
Priority: 7 7
SRLGs: not collected
Path Info:
  Incoming Address: 10.1.1.1
  Incoming:
  Explicit Route:
    No ERO

Route Exclusions:
  No XRO
Record Route: Disabled
Tspec: avg rate=4294967033 kbits, burst=1000 bytes, peak rate=4294967033 kbits
Session Attributes: Local Prot: Not Set, Node Prot: Not Set, BW Prot: Not Set
Resv Info: None
Record Route: Disabled
  Espec: avg rate=4294967033 kbits, burst=1000 bytes, peak rate=4294967033 kbits
Displayed 0 (of 0) heads, 0 (of 0) midpoints, 1 (of 1) tails
Displayed 0 up, 0 down, 0 recovering, 0 recovered heads

```

show rsvp neighbors

Displays information about RSVP neighbors.

```

Mon Jul 1 14:58:48.888 IST
Global Neighbor: 10.127.60.48
  Interface Neighbor  Interface
  -----
  10.127.60.48       MgmtEth0/RP0/CPU0/0

```

show lmp gmpls optical-uni

Verifies LMP configuration and state.

```

Mon Jul 1 14:55:35.492 IST

GMPLS Optical-UNI LMP Router ID: 10.105.57.101

LMP Neighbor
Name: ncslk, IP: 10.127.60.48, Owner: GMPLS Optical-UNI
LMP: Disabled
IPCC ID: 1, State Up
LMP UDP port: 701
Known via      : Configuration
Type           : Routed
Destination IP : 10.127.60.48
Source IP      : 10.105.57.101

```

Interface I/F	Lcl Interface ID	Lcl Link ID	Interface LMP state
Optics0/1/0/1	7	10.0.5.5	Up
Optics0/1/0/0	6	10.0.4.4	Up

```

Optics0/0/0/1          11          10.0.3.3          Up
Optics0/0/0/0          10          10.11.1.1         Up

```

General Troubleshooting

Collect and analyze the output of the following commands for any software issues.

- `show tech-support mpls traffic-eng file filename`
- `show tech-support mpls rsvp file filename`
- `show lmp clients`
- `show rsvp neighbors`
- `show mpls traffic-eng link-management optical-uni controller optics Rack/Slot/Instance/Port`
- `show mpls traffic-eng tunnels tunnel-id`

Problem	Solution
When NCS 2000 node cannot route the DWDM wavelength to the destination, it displays a generic error message as No Route to destination .	As a superuser, collect and analyze the diagnostic information by entering the following address at the browser. <code>http://ip-address-of-head-node/diagnostics/wson</code>

You May Be Also Interested In

- GMPLS UNI commands: [Cisco IOS XR MPLS Command Reference](#).
- [GMPLS Restoration and Reversion](#)

You May Be Also Interested In



CHAPTER 9

Remote Node Management

- [Remote node management using GCC](#) , on page 137
- [iBGP supports over GCC interfaces](#), on page 143


Remote node management using GCC

Remote node management using GCC is a network management method that

- leverages the General Communication Channel (GCC) embedded in optical transport networks,
- delivers reliable, out-of-band communication between centralized controllers and remote network nodes, and
- enables real-time monitoring, configuration, and maintenance activities without requiring direct physical access to each node.

From R7.2.1 onwards, the OTN-XP card provides OTU interface that supports communication channels between adjacent network elements or nodes using GCC bytes in the OTN header. Remote node management is supported over the GCC interface.

Figure 7: Remote Node Management in Linear Topology



The remote nodes can be dynamically discovered over the GCC interface using OSPF. The connectivity to the management network can be achieved using OSPF and static routes.



Note The GCC2 and GCC0 interfaces are supported in NCS 1004. The GCC0 interface is supported on the Coherent DSP controller whereas the GCC2 interface is supported on the ODU controller.



Note The GCC0 and GCC2 interfaces are supported in Muxponder and Muxponder slice modes. Only the GCC0 interface is supported in the Regeneration (Regen) mode.

From R7.2.1 onwards, the node supports GCC0 on the corresponding OTU2, OTU2e, and OTU4 interfaces. The node (Cisco FPGA) supports a maximum of 22 GCC channels for each card.



Note The GCC0 and GCC1 interfaces are supported on OTN-XP card and GCC2 interface is not supported.

Supported and unsupported features of the GCC interface

This table lists supported and unsupported features for remote node management using the GCC interface.

Table 12: Feature Support Overview:

Feature/Functionality	Supported on GCC Interface?	Notes
gRPC protocol	No	gRPC is not supported over the GCC interface.
Open Config	No	Not supported due to lack of gRPC support.
Streaming telemetry	No	Not supported due to lack of gRPC support.
Tx and Rx packet count statistics	Yes	Only Tx and Rx packet count information is available in GCC
Remote node management (after initial provisioning)	Yes	Devices can be managed over GCC only when connected through the management network using GCC.
Initial provisioning and bring-up via GCC	No	Must use console or management Ethernet interface for initial setup.
Remote management after headless/HA event	May be impacted	Events like reloads or driver restarts at intermediate nodes may affect management of subsequent nodes.
IP fragmentation for SCP protocol	No	Not supported; reduce packet size to less than 1454 bytes as a workaround.
TCP MSS configuration to avoid fragmentation	Yes	Use <code>tcp mss <maximum segment size></code> in global config mode.
Coherent DSP controller (QXP card) GCC0 interface	Yes	Supported on QXP card.
GCC0 speed on QXP card	Yes	7.7 Mbps.



Note For operations not supported on the GCC interface, use the console or management Ethernet interface as alternatives. For SCP protocol, configure TCP MSS or IPv4 MTU settings to avoid IP fragmentation issues.

Supported protocols

These protocols are supported over the GCC interface:

- PING
- SSH
- TELNET
- SCP
- TFTP
- FTP
- SFTP
- HTTP
- HTTPS
- OSPF

Enable the GCC interface

Use this task to enable GCC0, GCC1, or GCC2 interfaces on various line cards (1.2T, OTN-XP) to support management and communication channels.

Procedure

- Step 1** Enter configuration mode.
- Step 2** Configure the controller and the GCC interface for a line card.

If you want to configure	Then use the command
GCC2 interface on the 1.2T card	controller odu4 R/S/I/P/L gcc2
GCC0 interface on the 1.2T card	controller CoherentDSP R/S/I/P/L gcc2
GCC0 interface for the OTN-XP card	controller {otu2 otu2e otu4} R/S/I/P/L gcc0
GCC1 interface for the OTN-XP card	controller {odu2 odu2e odu4 oducn} R/S/I/P/L gcc1

Example:

This sample configuration enables the GCC2 interface for the 1.2T line card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller odu4 0/1/0/0/1
RP/0/RP0/CPU0:ios(config-odu)#gcc2
RP/0/RP0/CPU0:ios(config-odu)#commit
RP/0/RP0/CPU0:ios(config-odu)#exit
```

This sample configuration enables the GCC0 interface for the 1.2T line card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller CoherentDSP0/0/1/1
RP/0/RP0/CPU0:ios(config-odu)#gcc0
RP/0/RP0/CPU0:ios(config-odu)#commit
RP/0/RP0/CPU0:ios(config-odu)#exit
```

This sample configuration enables the GCC0 interface for the OTN-XP line card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller otu2 0/0/0/4/1
RP/0/RP0/CPU0:ios(config-odu)#gcc0
RP/0/RP0/CPU0:ios(config-odu)#commit
RP/0/RP0/CPU0:ios(config-odu)#exit
```

This sample configuration enables the GCC1 interface for the OTN-XP line card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#controller odu2 0/0/0/4/1
RP/0/RP0/CPU0:ios(config-odu)#gcc1
RP/0/RP0/CPU0:ios(config-odu)#commit
RP/0/RP0/CPU0:ios(config-odu)#exit
```

Configure the GCC interface

Use this task to configure the GCC0, GCC1, and GCC2 interfaces on 1.2T and OTN-XP cards using static or loopback IP addresses.

Procedure

- Step 1** Enter configuration mode.
- Step 2** Specify the GCC2 interface for a line card.

If you want to configure	Then use the command
GCC2 interface on the 1.2T card	interface gcc2 <i>R/S/I/P/L</i>
GCC0 interface on the 1.2T card	interface gcc0 <i>R/S/I/P</i>
GCC0 interface on the OTN-XP card	interface gcc0 <i>R/S/I/P</i>
GCC1 interface on the OTN-XP card	interface gcc1 <i>R/S/I/P</i>

- Step 3** Use the command **ipv4 address** *ipv4-address net-mask* to set the IPv4 address for the interface.

Example:

This sample configures the GCC2 interface using the static IP address on the 1.2T line card

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#interface gcc2 0/1/0/0/1
RP/0/RP0/CPU0:ios(config-if)#ipv4 address 198.51.100.244 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#commit
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show run interface gcc2 0/1/0/0/1
interface GCC20/1/0/0/1
ipv4 address 10.1.1.1 255.255.255.0
!
```

This sample configures the GCC2 interface using the loopback IP address on 1.2T card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:R2(config)#interface gcc2 0/1/0/0/1
RP/0/RP0/CPU0:R2(config-if)#ipv4 unnumbered loopback 0
RP/0/RP0/CPU0:ios(config-if)#exit
```

This sample checks the status of GCC2 interface.

```
RP/0/RP0/CPU0:ios#show ipv4 interface brief
Wed Sep 22 17:10:04.190 IST
Interface IP-Address Status Protocol Vrf-Name
GCC20/0/0/0/1 198.51.100.234 Up Up default
GCC20/3/0/1/3 198.51.100.244 Up Up default
Loopback0 198.51.100.224 Up
```

This sample configures the GCC0 interface using the static IP address on 1.2T or OTN-XP card. enables the GCC1 interface for the OTN-XP line card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#interface gcc0 0/1/0/0
P/0/RP0/CPU0:ios(config-if)#ipv4 address 198.51.100.244 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#commit
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
RP/0/RP0/CPU0:ios#show run interface gcc0 0/1/0/0
interface GCC00/1/0/0
ipv4 address 198.51.100.244 255.255.255.0
!
```

This sample configures the the GCC0 interface using the loopback IP address on 1.2T or OTN-XP card.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:R2(config)#interface gcc0 0/1/0/0
RP/0/RP0/CPU0:R2(config-if)#ipv4 unnumbered loopback 0
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
```

Configure static routes over the GCC interface

Use this task to configure the router to forward packets for specific networks or hosts via the GCC interface using manually defined routes.

Procedure

- Step 1** Enter configuration mode.
- Step 2** Enter the router static configuration mode.
- Step 3** Use the command **address-family ipv4 unicast** *ip4 address default-gateway* to enter address family configuration mode. This step also configures a routing session using standard IPv4 address prefixes.

Example:

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#router static address-family ipv4 unicast 0.0.0.0/0 10.105.57.1
RP/0/RP0/CPU0:ios(config)#exit
```

Configure OSPF routes over the GCC interface

Enable OSPF dynamic routing between gateway and remote nodes using GCC interfaces.

This task enables OSPF over GCC interfaces, which facilitates efficient OSPF communication in specialized network environments.

Procedure

	Command or Action	Purpose
Step 1	Enter configuration mode and enable the OSPF routing process using the command router ospf <i>process-id</i> router-id <i>ip-address</i> Example:	
Step 2	Assign the OSPF area and specify the interfaces to include using the command area <i>area-id</i> interface <i>type R/S/I/P/L</i>	
Step 3	On the remote node, redistribute connected routes into OSPF, using the command redistribute connected .	
Step 4	Exit the configuration mode upon completion. Example: Gateway Node: <pre>configure router ospf 1 router-id 192.0.2.89 area 0 interface Loopback0</pre>	

	Command or Action	Purpose
	<pre> ! interface MgmtEth0/RP0/CPU0/1 ! interface GCC20/0/0/0/1 ! interface GCC20/0/0/0/2 Remote Node: configure router ospf 1 router-id 192.0.2.92 redistribute connected area 0 interface Loopback0 ! interface GCC20/0/0/0/1 ! interface GCC20/0/0/0/2 </pre>	

OSPF is configured over GCC interfaces, enabling gateway and remote nodes to dynamically exchange routing information.

iBGP supports over GCC interfaces

iBGP support over GCC interfaces is a routing capability that

- allows external devices to exchange BGP routes through the management interfaces of NCS 1004 systems,
- enables NCS 1004 devices to advertise local networks and manage them using BGP-learned paths, and
- establishes iBGP sessions over GCC for exchanging BGP routes.

Caution: Follow restrictions for iBGP support using GCC

You can configure VRF on the GCC management interfaces (port 0 and port 1) of the NCS 1004 device to achieve traffic isolation between the two management ports.

NCS 1004 supports GCC0 and GCC2 interfaces for the 1.2T line card, allowing greater flexibility and connectivity options. The device advertises its local networks using BGP and manages them through paths learned from the iBGP sessions established over the GCC interfaces.

Restrictions for iBGP Support Using GCC

- IP fragmentation is not supported on the GCC interface.
- The BGP configuration over Open Config (OC) is not supported.



Note The limitations of Remote Node Management Using GCC are applicable for iBGP Support Using GCC. For more information, see [Limitations of Remote Node Management](#).

Caution: Follow restrictions for iBGP support using GCC

Consider these restrictions when implementing iBGP support using GCC:

- Do not use IP fragmentation on the GCC interface; this function is not supported.
- Do not configure BGP over Open Config (OC) on GCC interfaces; this configuration is not supported.
- Observe all limitations found for Remote Node Management using GCC; these also apply to iBGP support using GCC

Enabling the GCC Interface

To enable the GCC2 interface, use the following commands:

```
configure
controller odu4 R/S/I/P/L
gcc2
commit
exit
```

To enable the GCC0 interface, use the following commands:

```
configure
controller CoherentDSP R/S/I/P
gcc0
commit
exit
```

Configuring the Management Interface

To configure the management Ethernet interface with VRF, use the following commands:

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios#interface MgmtEth0/RP0/CPU0/1
RP/0/RP0/CPU0:ios(config-if)#vrf transport-vrf
```

```
RP/0/RP0/CPU0:ios(config-if)#ipv4 address ipv4-address  
RP/0/RP0/CPU0:ios(config-if)#commit  
RP/0/RP0/CPU0:ios(config-if)#exit
```

The following example displays how to configure the management Ethernet interface with VRF.

```
RP/0/RP0/CPU0:ios#configure  
RP/0/RP0/CPU0:ios#interface MgmtEth0/RP0/CPU0/1  
RP/0/RP0/CPU0:ios(config-if)#vrf transport-vrf  
RP/0/RP0/CPU0:ios(config-if)#ipv4 address 192.0.2.1 255.255.255.255  
RP/0/RP0/CPU0:ios(config-if)#commit  
RP/0/RP0/CPU0:ios(config-if)#exit
```

Configuring the Loopback Interface

To configure the loopback interface 0 with VRF, use the following commands:

```
RP/0/RP0/CPU0:ios#configure  
RP/0/RP0/CPU0:ios#interface Loopback0  
RP/0/RP0/CPU0:ios(config-if)#vrf transport-vrf  
RP/0/RP0/CPU0:ios(config-if)#ipv4 address ipv4-address  
RP/0/RP0/CPU0:ios(config-if)#commit  
RP/0/RP0/CPU0:ios(config-if)#exit
```

The following example displays how to configure the loopback interface 0 with VRF.

```
RP/0/RP0/CPU0:ios#configure  
RP/0/RP0/CPU0:ios#interface Loopback0  
RP/0/RP0/CPU0:ios(config-if)#vrf transport-vrf  
RP/0/RP0/CPU0:ios(config-if)#ipv4 address 192.0.2.1 255.255.255.255  
RP/0/RP0/CPU0:ios(config-if)#commit  
RP/0/RP0/CPU0:ios(config-if)#exit
```

Configuring the GCC interface

To configure the GCC2 interface with VRF and static IP address, use the following commands:

```
configure  
interface gcc2 R/S/I/P/L  
vrf transport-vrf  
ipv4 address ipv4-address  
commit  
exit
```

To configure the GCC0 interface with VRF and static IP address, use the following commands:

```
configure  
interface gcc0 R/S/I/P  
vrf transport-vrf  
ipv4 address ipv4-address
```

commit

exit

Examples

The following sample displays how to configure the GCC2 interface with VRF and static IP address.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#interface gcc2 0/1/0/0/1
RP/0/RP0/CPU0:ios(config-if)#vrf transport-vrf
P/0/RP0/CPU0:ios(config-if)#ipv4 address 198.51.100.5 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#commit
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
```

The following sample displays how to configure the GCC2 interface using loopback IP address.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:R2(config)#interface gcc2 0/1/0/0/1
RP/0/RP0/CPU0:R2(config-if)#ipv4 unnumbered loopback 0
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
```

The following sample displays how to configure the GCC0 interface with VRF and static IP address.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:ios(config)#interface gcc0 0/1/0/0
RP/0/RP0/CPU0:ios(config-if)#vrf transport-vrf
P/0/RP0/CPU0:ios(config-if)#ipv4 address 198.51.100.2 255.255.255.0
RP/0/RP0/CPU0:ios(config-if)#commit
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
```

The following sample displays how to configure the GCC0 interface using the loopback IP address.

```
RP/0/RP0/CPU0:ios#configure
RP/0/RP0/CPU0:R2(config)#interface gcc0 0/1/0/0
RP/0/RP0/CPU0:R2(config-if)#ipv4 unnumbered loopback 0
RP/0/RP0/CPU0:ios(config-if)#exit
RP/0/RP0/CPU0:ios(config)#exit
```

Verifying iBGP Support Using GCC

To verify BGP support using GCC configuration, use the following **show** commands:

```
RP/0/RP0/CPU0:ios#show bgp vrf transport-vrf neighbors brief
Neighbor      Spk    AS Description                               Up/Down  NBRState
198.51.100.0    0     200                                           00:51:49 Established
198.51.100.1    0     100                                           00:50:32 Established
```

```
RP/0/RP0/CPU0:ios#show bgp vrf transport-vrf
BGP VRF transport-vrf, state: Active
BGP Route Distinguisher: 192.0.2.7:0
VRF ID: 0x60000002
BGP router identifier 192.0.2.7, local AS number 100
Non-stop routing is enabled
BGP table state: Active
```

```

Table ID: 0xe0000002   RD version: 51
BGP main routing table version 51
BGP NSR Initial initsync version 11 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 192.0.2.7:0 (default for vrf transport-vrf)

*> 209.165.201.30/27      198.51.100.0          0          0 200 i
*> 209.165.201.28/27      0.0.0.0                0          32768 i
*> 209.165.201.26/27      0      100          0 i
*> 209.165.201.24/27      198.51.100.2          0      100          0 300 i

```

```

RP/0/RP0/CPU0:ios#show bgp vrf transport-vrf
BGP VRF transport-vrf, state: Active
BGP Route Distinguisher: 203.0.113.10:0
VRF ID: 0x60000002
BGP router identifier 203.0.113.10, local AS number 100
Non-stop routing is enabled
BGP table state: Active
Table ID: 0xe0000002   RD version: 51
BGP main routing table version 51
BGP NSR Initial initsync version 11 (Reached)
BGP NSR/ISSU Sync-Group versions 0/0

Status codes: s suppressed, d damped, h history, * valid, > best
                i - internal, r RIB-failure, S stale, N Nexthop-discard
Origin codes: i - IGP, e - EGP, ? - incomplete
      Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 203.0.113.10:0 (default for vrf transport-vrf)

*> 209.165.201.30/27      198.51.100.0          0          0 200 i
*> 209.165.201.28/27      0.0.0.0                0          32768 i
*>i209.165.201.26/27      198.51.100.12         0      100          0 i
*>i209.165.201.24/27      198.51.100.24         0      100          0 300 i


```

iBGP configuration parameters for GCC interfaces

This table summarizes the configuration required for enabling iBGP sessions between two NCS 1004 devices.

Use case:

Consider two NCS 1004 devices, R2 and R3, directly connected through GCC0 interfaces. In this use case, R2 (with IP address 198.51.100.21) and R3 (with IP address 198.51.100.22) are configured with the commands to establish an iBGP session using their respective transport-VRFs.



R2 is connected through GCC0 0/0/0/0 interface with IP address of 198.51.100.21 and R3 is connected through GCC0 0/1/0/0 with IP address of 198.51.100.22. The R2 and R3 devices are connected to external devices through management interfaces.

Table 13: iBGP configuration commands for NCS 1004

Configuration on R2	Configuration on R3
<p>Global Configuration on R2</p> <pre>hw-module location 0/0 mxponder trunk-rate 600G client-rate 100GE vrf transport-vrf address-family ipv4 unicast</pre>	<p>Global Configuration on R3</p> <pre>hw-module location 0/0 mxponder trunk-rate 600G client-rate 100GE vrf transport-vrf address-family ipv4 unicast</pre>
<p>Interface Configuration on R2</p> <pre>interface Loopback0 vrf transport-vrf ipv4 address 192.0.2.2 255.255.255.255 interface MgmtEth0/RP0/CPU0/1 vrf transport-vrf ipv4 address 198.51.100.25 255.255.255.0 controller ODU40/0/0/0/2 gcc2 interface GCC20/0/0/0/2 vrf transport-vrf ipv4 address 198.51.100.21 255.255.255.0</pre>	<p>Interface Configuration on R3</p> <pre>interface Loopback0 vrf transport-vrf ipv4 address 203.0.113.3 255.255.255.255 interface MgmtEth0/RP0/CPU0/1 vrf transport-vrf ipv4 address 198.51.100.32 255.255.255.0 controller ODU40/1/0/0/2 gcc2 interface GCC20/1/0/0/2 vrf transport-vrf ipv4 address 198.51.100.22 255.255.255.0</pre>
<p>Route-policy Configuration on R2</p> <pre>route-policy PASS-ALL pass end-policy</pre>	<p>Router Policy Configuration on R3</p> <pre>route-policy PASS-ALL pass end-policy</pre>
<p>Static Route Configuration on R2</p> <pre>router static address-family ipv4 unicast 0.0.0.0/0 198.51.100.28 ! vrf transport-vrf address-family ipv4 unicast 198.51.100.0/24 198.51.100.22</pre>	<p>Static Route Configuration on R3</p> <pre>router static address-family ipv4 unicast 0.0.0.0/0 198.51.100.28 ! vrf transport-vrf address-family ipv4 unicast 198.51.100.0/24 198.51.100.21</pre>

Configuration on R2	Configuration on R3
<p>BGP Configuration on R2</p> <pre> router bgp 100 bgp router-id 192.0.2.123 address-family vpnv4 unicast ! vrf transport-vrf rd auto address-family ipv4 unicast network 203.0.113.1/32 ! neighbor 198.51.100.22 remote-as 100 address-family ipv4 unicast route-policy PASS-ALL in route-policy PASS-ALL out next-hop-self ! </pre>	<p>BGP Configuration on R3</p> <pre> router bgp 100 bgp router-id 192.0.2.124 address-family vpnv4 unicast ! vrf transport-vrf rd auto address-family ipv4 unicast network 203.0.113.3/32 ! neighbor 198.51.100.21 remote-as 100 address-family ipv4 unicast route-policy PASS-ALL in route-policy PASS-ALL out next-hop-self ! </pre>
<p>BGP Verification on R2</p> <pre> RP/0/RP0/CPU0:ios#show bgp sessions Mon Jul 20 14:47:30.378 UTC Neighbor VRF Spk AS InQ OutQ NBRState NSRState 198.51.100.22 transport-vrf 0 100 0 0 Established None </pre>	<p>BGP Verification on R3</p> <pre> RP/0/RP0/CPU0:regen#show bgp sessions Tue Jul 21 02:50:14.134 UTC Neighbor VRF Spk AS InQ OutQ NBRState NSRState 198.51.100.21 transport-vrf 0 100 0 0 Established None </pre>



CHAPTER 10

Smart Licensing

This chapter describes the smart licensing configuration on Cisco NCS 1004.

- [Understanding Smart Licensing, on page 151](#)
- [Configure Smart Licensing, on page 155](#)
- [Smart Licensing for OTN-XP Line Card, on page 161](#)

Understanding Smart Licensing

Smart Licensing is a cloud-based approach to licensing. Smart Licensing simplifies the licensing experience across the enterprise making it easier to purchase, deploy, track, and renew Cisco Software. It provides visibility into license ownership and consumption through a single, simple user interface. The solution allows you to easily track the status of your license and software usage trends.

Smart Licensing helps you simplify three core functions:

- **Purchasing:** The software that you have installed in your network can be registered, without Product Activation Keys (PAKs).
- **Management:** You can automatically track activations against your license entitlements. Also, there is no need to install the license file on every node. You can create license pools (logical grouping of licenses) to reflect your organization structure. Smart Licensing offers you Cisco Smart Software Manager, a centralized portal that enables you to manage all your Cisco software licenses from one centralized website.
- **Reporting:** Through the portal, Smart Licensing offers an integrated view of the licenses you have purchased and what has been deployed in your network. You can use this data to make better purchasing decisions, based on your consumption.

Smart Licensing Features

- Your device initiates a call home and requests the licenses it needs.
- Pooled licences - Licences are company account-specific, and can be used with any compatible device in your company. You can activate or deactivate different types of licenses on the device without actually installing a license file on the device.
- Licenses are stored securely on Cisco servers.
- Licenses can be moved between product instances without license transfer. This greatly simplifies the reassignment of a software license as part of the Return Material Authorization (RMA) process.

- It provides a complete view of all the Smart Software Licenses used in the network using a consolidated usage report of software licenses and devices in one easy-to-use portal.

Cisco Smart Account

NCS 1004 integrates with Cisco Smart Accounts to simplify license management.

Smart Accounts provide a centralized, organized, and simple-to-use solution for managing Cisco software licenses across an entire organization. They act as a container that holds all Cisco software assets, allowing customers to view, store, manage, and move these assets as needed. Smart Accounts enable full visibility into license entitlements and usage, helping optimize software management. Smart Accounts support license pooling, portability, and provide compliance reporting.

When creating a Smart Account, you must have the authority to represent the requesting organization. After you submit the request, it goes through a brief approval process. Access <http://software.cisco.com> to learn about, set up, or manage Smart Accounts.

Cisco Smart Software Manager Overview

Cisco Smart Software Manager enables you to manage all your Cisco Smart software licenses from one centralized website. With Cisco Smart Software Manager, you organize and view your licenses in groups called virtual accounts (collections of licenses and product instances). Use the Cisco Smart Software Manager to do the following tasks:

- Create, manage, or view virtual accounts.
- Create and manage Product Instance Registration Tokens.
- Transfer licenses between virtual accounts or view licenses.
- Transfer, remove, or view product instances.
- Run reports against your virtual accounts.
- Modify your email notification settings.
- View overall account information.

Virtual Accounts

Virtual Accounts are customizable subaccounts within a Smart Account used to organize and optimize Cisco licenses. They can be structured to reflect business units, product types, or geographic locations, allowing better planning and utilization of assets. They are created and maintained by the Smart Account administrator. Smart Licencing allows you to create multiple license pools or virtual accounts within the Smart Software Manager portal. Using the Virtual Accounts option that you can aggregate licenses into discrete bundles that are associated with a cost center so that one section of an organization cannot use the licenses of another section of the organization. For example, if you segregate your company into different geographic regions, you can create a virtual account for each region to hold the licenses and product instances for that region.

All new licenses and product instances are placed in the default virtual account in the Smart Software Manager, unless you specify a different one during the order process. After you access the default account, you may choose to transfer them to any other account, provided you have the required access permissions.

Use the Smart Software Manager portal to create license pools or transfer licenses.

Product Instance Registration Tokens

A product requires a registration token until you have registered the product. On successful registration, the device receives an identity certificate. This certificate is saved and automatically used for all future

communications with Cisco. Registration tokens are stored in the Product Instance Registration Token Table that is associated with your enterprise account. Registration tokens can be valid 1–365 days.


Product Instances

A product instance is an individual device with a unique device identifier (UDI) that is registered using a product instance registration token (or registration token). You can register any number of instances of a product with a single registration token. Each product instance can have one or more licenses residing in the same virtual account. Product instances must periodically connect to the Cisco Smart Software Manager servers during a specific renewal period. If you remove the product instance, its licenses are released and made available within the virtual account.

Smart Licensing Work Flow

The following figure depicts a working model of smart licensing that involves a three-step procedure.

Figure 8: Smart Licensing Work Flow



1. **Setting up Smart Licensing:** You can place the order for Smart Licensing, to manage licenses on the Cisco.com portal. You agree to the terms and conditions governing the use and access of Smart Licensing in the Smart Software Manager portal.
2. **Enabling and Use Smart Licensing:** Smart Licensing is enabled by default. You can use either of the following options to communicate:
 - **Smart Call Home:** The Smart Call Home feature is automatically configured when Smart Licensing is enabled. Smart Call Home is used by Smart Licensing as a medium for communication with the Cisco license service. Call Home feature allows Cisco products to periodically call-home and perform an audit and reconciliation of your software usage information. This information helps Cisco efficiently track your install base, keep them up and running, and effectively pursue service and support contract renewals. For more information on Smart Call Home feature, see http://www.cisco.com/c/dam/en/us/td/docs/switches/lan/smart_call_home/SCH_Deployment_Guide.pdf.
 - **Smart Software Manager Satellite :** is a component of Cisco Smart Licensing and works with Cisco Smart Software Manager (SSM). It helps customers intelligently manage product licenses, providing near real-time visibility and reporting of the Cisco licenses they purchase and consume. For customers who do not want to manage their installed base using a direct Internet connection, the Smart Software Manager satellite is installed on the customer premises and provides a subset of

Cisco SSM functionality. After you download the satellite application, deploy it, and register it to Cisco SSM, you can perform the following functions locally:

- Activate or register a license
- Get visibility to your company's licenses
- Transfer licenses between company entities

Periodically, the satellite must synchronize with Cisco SSM to reflect the latest license entitlements.

For more information about the Smart Software Manager satellite, see <http://www.cisco.com/c/en/us/buy/smart-accounts/software-manager-satellite.html>.

3. **Manage and Report Licenses:** You can manage and view reports about your overall software usage in the Smart Software Manager portal. Compliance reporting describes the types of Smart Licensing reports.

Benefits of Smart Licensing

- Licenses are not locked to perform configurations even if the license limit exceeds the paid license limit. You are notified with out-of-compliance notification to buy additional licenses when the license limit exceeds the paid license limit. This saves time with the ability to transfer licenses across the organization.
- Licenses can be pooled across the entire organization, enabling them to be reused across organizational boundaries.
- Provides software asset management information so that you can plan and track the licenses.

Licensing in NCS 1004

Cisco NCS 1004 has the following line card PIDs :

- **NCS1K4-1.2T-K9**—High-cost PID. You can use this line card without any explicit licensing.
- **NCS1K4-1.2T-L-K9**—Licensed PID for 1.2T line card and the licenses are charged per port.
- **NCS1K4-OTN-XPL**—Licensed PID for OTN-XP line card.
- **NCS1K4-QXP-L-K9**—Licensed PID for QXP line card.
- **NCS1K4-QXP-K9** —Non-licensed PID for QXP line card.

You can use either one or a combination of both types of the 1.2T line card in the NCS 1004.

Software Entitlements of Cisco NCS 1004

Software entitlement is a system that consists of a license manager on Cisco NCS 1004. The license manager manages licenses for various software and hardware features. The license manager parses and authenticates the license before accepting it.

The following table lists the features and its corresponding entitlements that can be enabled on Cisco NCS 1004 using licenses:

Table 14: Software Entitlements of Cisco NCS 1004

Feature	Software Entitlement
NCS1K4 Smart License-one QSFP28 client	S-NCS1K4-LIC-100G=
NCS1K4 Smart License - one QSFP28 client with encryption	S-NCS1K4-LIC-100X=
NCS1K4 Smart License - 100Gbps of client bandwidth	S-NCS1K4-100G-CL=

The licenses are charged per port basis and dependent on the number of trunk ports and client ports that you configure. The license count for the configuration of 4 x 100GE client ports or lesser is zero. For configurations greater than 4 x 100GE client ports, the license count is incremented by one for every 100GE client port configured at the slice level. The license count for the trunk port is incremented based on the BPS & optics configuration.

Configure Smart Licensing

To configure smart licensing in Cisco NCS 1004, perform the following tasks:

Procedure

Step 1 Configure the domain name server for the smart license server.

Example:

```
RP/0/RP0/CPU0:ios#configure
Sat Dec 15 15:25:14.385 IST
RP/0/RP0/CPU0:NCS1004(config)#domain name-server 198.51.100.247
```

Step 2 Setup the CiscoTAC-1 profile and destination address for Smart Call Home, using the following commands:

call-home

service active

contact smart-licensing

profile CiscoTAC-1

active

destination address http {http|https}://{FQDN}/its/service/oddce/services/DDCEService

destination transport-method http

Note

FQDN must be either Cisco Smart Software Manager FQDN (tools.cisco.com) or Smart Licensing satellite server FQDN. You must configure the DNS server before setting-up the call-home destination address as FQDN. Use the **domain name-server {DNS server IP}** command to configure the DNS server on the device.

Example:

```
domain name-server 198.51.100.247
call-home
```

```

service active
contact smart-licensing
profile CiscoTAC-1
active
destination address http https://tools.cisco.com/its/service/oddce/services/DDCEService
destination transport-method http

```

Note

CiscoTAC-1 profile is the default profile for smart licensing and it must not be deleted.

- Step 3** Configure the crypto ca Trust point profile, if CRL distribution point is not defined in the Satellite server certificate or if the device is not able to reach the host mentioned in the CRL distribution point.

Example:

```
RP/0/RP0/CPU0:ios(config)#crypto ca trustpoint Trustpool crl optional
```

- Step 4** Create and copy the registration token ID using Cisco Smart Software Manager.

For more details about creating a token, see [Creating a Token, on page 156](#).

- Step 5** In the privileged EXEC mode, register the token ID in Cisco NCS 1004, using the following commands:

license smart register idtoken *token-ID*

The registration may fail if the token is invalid or there is communication failure between the device and the portal or satellite. If there is a communication failure, there is a wait time of 24 hours before the device attempts to register again. To force the registration, use the **license smart register idtoken** *token-ID* **force** command.

When your device is taken off the inventory, shipped elsewhere for redeployment, or returned to Cisco for replacement using the return merchandise authorization (RMA) process, you can use the **license smart deregister** command to cancel the registration on your device. All smart licensing entitlements and certificates on the platform are removed.

ID certificates are renewed automatically after six months. In case, the renewal fails, the product instance goes into unidentified state. You can manually renew the ID certificate using the **license smart renew id** command.

Authorization periods are renewed by the Smart Licensing system every 30 days. As long as the license is in an 'Authorized' or 'Out-of-Compliance' (OOC), the authorization period is renewed. Use the **license smart renew auth** command to make an on-demand manual update of your registration. Thus, instead of waiting 30 days for the next registration renewal cycle, you can issue this command to instantly find out the status of your license.

After 90 days, the authorization period expires and the status of the associated licenses display "AUTH EXPIRED". Use the **license smart renew auth** command to retry the authorization period renewal. If the retry is successful, a new authorization period begins.

Creating a Token

To create a new token using Cisco Smart Software Manager, perform the following tasks:

Procedure

-
- Step 1** Log in to the Cisco Smart Software Manager.
URL: <https://software.cisco.com/#SmartLicensing-Inventory>
- Step 2** Select the appropriate Virtual Account.
- Step 3** From the **General** tab, choose **New Token**. Create Registration Token dialog box appears.
- Step 4** Follow the dialog to provide a name, duration, and export compliance applicability before accepting the terms and responsibilities.
- Step 5** Click **Create Token**.
-

Verifying Smart Licensing Configuration

After enabling Smart Licensing, you can use the **show** commands to verify the default Smart Licensing configuration. If any issue is detected, take corrective action before making further configurations.

- **show license all**
- **show license trace all**
- **show license status**
- **show license summary**
- **show license tech**
- **Show license udi**
- **show license usage**
- **show license platform detail**
- **show license platform summary**
- **show license platform trace**
- **Show license platform trace all**
- **show tech-support smartlic**
- **show call-home detail**
- **show call-home trace all**
- **show tech-support call-home**

The following table defines the available license authorization status in Cisco NCS 1004:

Table 15: License Authorization Status

License Authorization Status	Description
Unconfigured	Smart Software Licensing is not configured.
Unidentified	Smart Software Licensing is enabled but is not registered.
Registered	Device registration is completed and an ID certificate is received that is used for future communication with the Cisco licensing authority.
Authorized	Registration is completed with a valid Smart Account and license consumption has begun. This indicates compliance.
Out of Compliance	Consumption exceeds available licenses in the Smart Account.
Authorization Expired	The device is unable to communicate with the Cisco Smart Software Manager for an extended period. This state occurs after 90 days of expiry. The device attempts to contact the CSSM every hour to renew the authorization until the registration period expires.

Example 1:

The following example shows the sample output of the **show license all** command.

```
RP/0/RP0/CPU0:ios#show license all
Mon Feb 11 15:58:44.047 IST

Smart Licensing Status
=====

Smart Licensing is ENABLED

Registration:
  Status: REGISTERED
  Smart Account: BU Production Test
  Virtual Account: NCS1000
  Initial Registration: SUCCEEDED on Mon Feb 11 2019 15:51:10 IST
  Last Renewal Attempt: None
  Next Renewal Attempt: Sat Aug 10 2019 15:52:10 IST
  Registration Expires: Tue Feb 11 2020 15:46:59 IST

License Authorization:
  Status: AUTHORIZED on Mon Feb 11 2019 15:53:40 IST
  Last Communication Attempt: SUCCEEDED on Mon Feb 11 2019 15:53:40 IST
  Next Communication Attempt: Wed Mar 13 2019 15:53:39 IST
  Communication Deadline: Sun May 12 2019 15:47:29 IST

License Usage
=====

NCS1K4 smart license - one QSFP28 client (S-NCS1K4-LIC-100G=):
  Description: NCS1K4 smart license - one QSFP28 client
  Count: 8
  Version: 1.0
  Status: AUTHORIZED

NCS1K4 smart license - one QSFP28 client with encryption (S-NCS1K4-LIC-100X=):
  Description: NCS1K4 smart license - one QSFP28 client with encryption
  Count: 8
```

```

Version: 1.0
Status: AUTHORIZED

Product Information
=====
UDI: SN:CAT2231B18Y,UUID:default-sdr

Agent Version
=====
Smart Agent for Licensing: 2.2.0_rel/48

```

Example 2:

The following example shows the sample output of the **show license platform detail** command.

```

RP/0/RP0/CPU0:ios#show license platform detail
Mon Feb 11 15:59:55.422 IST
Current state:    REGISTERED

Collection: LAST: Mon Feb 11 2019 15:57:53 IST
              NEXT: Mon Feb 11 2019 16:57:53 IST
Reporting:  LAST: Mon Feb 11 2019 15:57:53 IST
              NEXT: Tue Feb 12 2019 15:57:53 IST

Parameters: Collection interval:      60 minute(s)
              Reporting interval:    1440 minute(s)
              Throughput gauge:     1000000 Kbps

=====
Feature/Area 'system'
  Name: System
  Status: ACTIVE
  Flags: CONFIG

  [ 1] Name: NCS1K4 smart license - one QSFP28 client
        Entitlement Tag:
regid.2018-05.com.cisco.S-NCS1K4-LIC-100G=,1.0_03df009f-5ac5-48da-af50-4279ddea5e24
        Count: Last reported:      8
              Next report:        0
  [ 2] Name: NCS1K4 smart license - one QSFP28 client with encryption
        Entitlement Tag:
regid.2018-05.com.cisco.S-NCS1K4-LIC-100X=,1.0_3938b0c5-f635-4426-9f0f-936d930cea9e
        Count: Last reported:      8
              Next report:        0

```

Example 3:

The following example shows the sample output of the **show license status** command.

```

RP/0/RP0/CPU0:ios#show license status
Mon Feb 11 16:02:24.499 IST

Smart Licensing is ENABLED
  Initial Registration: SUCCEEDED on Mon Feb 11 2019 15:51:10 IST
  Last Renewal Attempt: None
  Next Renewal Attempt: Sat Aug 10 2019 15:52:10 IST
  Registration Expires: Tue Feb 11 2020 15:46:59 IST

License Authorization:
  Status: AUTHORIZED on Mon Feb 11 2019 15:53:40 IST
  Last Communication Attempt: SUCCEEDED on Mon Feb 11 2019 15:53:40 IST
  Next Communication Attempt: Wed Mar 13 2019 15:53:39 IST

```

Communication Deadline: Sun May 12 2019 15:47:29 IST

Example 4:

The following example shows the sample output of the **show license usage** command.

```
RP/0/RP0/CPU0:ios#show license usage

Mon Feb 11 15:59:29.817 IST

License Authorization:
  Status: AUTHORIZED on Mon Feb 11 2019 15:53:40 IST

NCS1K4 smart license - one QSFP28 client (S-NCS1K4-LIC-100G=):
  Description: NCS1K4 smart license - one QSFP28 client
  Count: 8
  Version: 1.0
  Status: AUTHORIZED

NCS1K4 smart license - one QSFP28 client with encryption (S-NCS1K4-LIC-100X=):
  Description: NCS1K4 smart license - one QSFP28 client with encryption
  Count: 8
  Version: 1.0
  Status: AUTHORIZED
```

Example 5:

The following example shows the sample output of the **show license udi** command.

```
RP/0/RP0/CPU0:ios#show license udiMon Feb 11 16:02:46.733 IST

Product Information
=====
UDI: SN:CAT2231B18Y,UUID:default-sdr
```

License Registration

You can use the following procedure to register license:

Procedure

-
- Step 1** Register the license using the following command:

```
RP/0/RP0/CPU0:ios#license smart register idtoken <idtoken>
```
 - Step 2** Browse to the URL : <https://software.cisco.com/software/cs/ws/platform/home#module/SmartLicensing>.
 - Step 3** Click **Inventory**.
 - Step 4** Click **Product Instances**.
 - Step 5** Select the node instance.
 - Step 6** Click **Actions**.
 - Step 7** Click **Remove**.
 - Step 8** Renew the authorization period using the following command:

```
RP/0/RP0/CPU0:ios#license smart renew auth
```

```
RP/0/RP0/CPU0:ios#show logging | i "Data and signature"
Thu May 27 09:57:02.237 UTC
RP/0/RP0/CPU0:May 27 09:54:57.783 UTC: smartlicserver[311]:
LICENSE-SMART_LIC-3-AUTH_RENEW_FAILED : Authorization renewal with the Cisco Smart Software
  Manager (CSSM) :
Error received from Smart Software Manager: Data and signature do not match for udi
PID:8812,SN:FOX2202WIVM
```

Note

The error message in the output of the **show logging** command is expected and is due to loss of synchronization between the CSSM server and the device after removing the product instance directly from the CSSM server.

Step 9 Perform deregister using the following command:

```
RP/0/RP0/CPU0:ios#license smart deregister

RP/0/RP0/CPU0:ios#show logging | i DEREG
Thu May 27 14:48:58.170 UTC
RP/0/RP0/CPU0:May 27 09:58:58.464 UTC: smartlicserver[311]:
%LICENSE-SMART_LIC-3-AGENT_DEREG_FAILED : Smart Agent for Licensing DeRegistration with
Cisco Smart Software Manager (CSSM) failed:
Agent received a failure status in a response message. Please check the Agent log file for
the detailed message.
```

Note

The error message in the output of the **show logging** command is expected.

Smart Licensing for OTN-XP Line Card

Overview

- The license calculation is based on 100G client bandwidth and is independent of the client type.
- The licensed OTN-XP Line Card PID is NCS1K4-OTN-XPL.
- The license is charged based on the usage of 100G client bandwidth.

Checking the License Usage Count

You can also check the number of licenses utilised, by entering the **show license all** command.

Configuring Slice

The following sample shows the configuration of slice 0 in Muxponder mode.

```
RP/0/RP0/CPU0:ios(config)#hw-module location 0/1 mxponder-slice 0
  trunk-rate 100G
  client-port-rate 2 lane 3 client-type OTU2
  client-port-rate 2 lane 4 client-type OTU2E
  client-port-rate 4 lane 1 client-type 10GE
  client-port-rate 4 lane 2 client-type OTU2
  client-port-rate 4 lane 3 client-type OTU2E
  client-port-rate 4 lane 4 client-type 10GE
  client-port-rate 5 lane 1 client-type OTU2E
  client-port-rate 5 lane 2 client-type 10GE
```

```
client-port-rate 5 lane 3 client-type OTU2
client-port-rate 5 lane 4 client-type OTU2E
```

Checking the Slice State

The following sample shows the status of the configured slice as **Provisioned**.

```
RP/0/RP0/CPU0:ios#show hw-module location 0/1 mxponder
Fri Dec 6 02:50:32.858 UTC
```

```
Location:          0/1
Slice ID:          0
Client Bitrate:    MIXED
Trunk Bitrate:     100G
Status:        Provisioned
LLDP Drop Enabled: FALSE
Client Port
```

Mapper/Trunk Port Traffic Split Percentage	Peer/Trunk Port	OTU40/1/0/0
OTU20/1/0/2/3	NONE	ODU20/1/0/0/2/3 100
OTU20/1/0/4/2	NONE	ODU20/1/0/0/4/2 100
OTU20/1/0/5/3	NONE	ODU20/1/0/0/5/3 100
OTU2E0/1/0/2/4	NONE	ODU2E0/1/0/0/2/4 100
OTU2E0/1/0/4/3	NONE	ODU2E0/1/0/0/4/3 100
OTU2E0/1/0/5/1	NONE	ODU2E0/1/0/0/5/1 100
OTU2E0/1/0/5/4	NONE	ODU2E0/1/0/0/5/4 100
TenGigEctr1r0/1/0/4/1	ODU2E0/1/0/0/4/1	NONE 100
TenGigEctr1r0/1/0/4/4	ODU2E0/1/0/0/4/4	NONE 100
TenGigEctr1r0/1/0/5/2	ODU2E0/1/0/0/5/2	NONE 100

Checking the License Count

The following sample shows the license usage count as 1.

```
RP/0/RP0/CPU0:ios#show license all
Fri Dec 6 02:58:39.906 UTC
```

```
Smart Licensing Status
=====
```

```
Smart Licensing is ENABLED
```

Registration:

```
Status: REGISTERED
Smart Account: BU Production Test
Virtual Account: NCS1000
Export-Controlled Functionality: ALLOWED
Initial Registration: SUCCEEDED on Dec 06 2019 02:54:27 UTC
Last Renewal Attempt: None
Next Renewal Attempt: Jun 03 2020 02:54:26 UTC
Registration Expires: Dec 05 2020 02:49:42 UTC
```

License Authorization:

```
Status: AUTHORIZED on Dec 06 2019 02:56:50 UTC
Last Communication Attempt: SUCCEEDED on Dec 06 2019 02:56:50 UTC
Next Communication Attempt: Dec 06 2019 14:56:49 UTC
Communication Deadline: Mar 05 2020 02:52:06 UTC
```

Export Authorization Key:

```
Features Authorized:
<none>
```

Utility:

```
Status: DISABLED
Data Privacy:
Sending Hostname: yes
```

```
Callhome hostname privacy: DISABLED
Smart Licensing hostname privacy: DISABLED
Version privacy: DISABLED
```

```
Transport:
  Type: Callhome
```

```
License Usage
=====
```

```
NCS 1K4 smart License - 100Gbps of client bandwidth (S-NCS1K4-100G-CL=):
  Description: 100G Client bandwidth - Smart License
  Count: 1
  Version: 1.0
  Status: Authorized
  Export status: NOT RESTRICTED
```

```
Product Information
=====
UDI: SN:CAT2217B025,UUID:default-sdr
```

```
Agent Version
=====
Smart Agent for Licensing: 4.10.4_rel/21
```

```
Reservation Info
=====
License reservation: DISABLED
```



Note The license count for 10GE, OTU2, and OTU2e is calculated as follows:

- If $1 \leq \text{number of client ports used} \leq 10$, it implies the "S-NCS1K4-100G-CL=" has license count of 1 and similarly if $11 \leq \text{number of client ports used} \leq 20$, it implies the "S-NCS1K4-100G-CL=" has license count of 2 and so on.
-



CHAPTER 11

USB Device Automount

- [USB automounts, on page 165](#)
- [Mount a USB device in virtual machine environments, on page 165](#)
- [Unmount a USB device, on page 166](#)

USB automounts

A USB automount is a system feature that

- automatically mounts a USB device in the sysadmin-vm with read and write permissions,
- supports automount only when the USB device file system is FAT or FAT32, and
- allows you to access files and folders on the mounted USB device through the disk2: file system, without needing to explicitly mount the device.

Operational details

- When you insert a USB device into NCS 1004, it is automatically mounted in sysadmin-vm with read and write permissions, and, by default, is unmounted in XR-vm.
- If the USB file system is not FAT or FAT32, the device will not be automounted in sysadmin-vm.
- To access or transfer files, you can use the disk2: file system reference.
- You must unmount the USB device from sysadmin-vm before physically removing it from NCS 1004.
- If you need to remount the USB device after unmounting but before removing it, use the `mount` command.
- You can mount the USB device in both XR-vm and sysadmin-vm at the same time. Before physically removing the USB device from NCS 1004, unmount it from both sysadmin-vm and XR-vm

Mount a USB device in virtual machine environments

Use this task to mount a USB device so you can access, copy, and manage files within your virtual machine environment.

Mounting a USB device integrates external storage into your virtual machine, allowing you to move essential files between systems or for backup purposes. Procedures differ slightly between sysadmin-vm and XR-vm.

Procedure

Step 1 Use the command **usb device operation mount** to mount the USB device in sysadmin-vm.

Example:

```
sysadmin-vm:0_RP0#usb device operation mount
Fri Jul 13 09:26:00.821 UTC success usb mounted
```

Step 2 Use the command **unmount disk2: undo** to mount the USB device in XR-vm.

Example:

```
RP/0/RP0/CPU0:ios#unmount disk2: undo
Fri Jul 13 14:56:34.326 IST
disk2: mounted successfully.
```

Step 3 Use the command **scp** to copy the file to the USB device.

Example:

```
[sysadmin-vm:0_RP0:~/showtech]$scp showtech-envmon-admin-2018-Jul-04.171400.IST.tgz /disk2\:
[sysadmin-vm:0_RP0:~/showtech]$cd /disk2\:
[sysadmin-vm:0_RP0:/disk2:]$ls -lrt
total 122424
drwxr-xr-x 2 root root      8192 Jul 12  2017 System Volume Information
drwxr-xr-x 2 root root      8192 Jun 11 16:16 boot
drwxr-xr-x 3 root root      8192 Jun 11 16:17 EFI
-rwxr-xr-x 1 root root 125306880 Jul 10 13:50 calvVarLog.tar
-rwxr-xr-x 1 root root   23023 Jul 13 05:23 showtech-envmon-admin-2018-Jul-04.171400.IST.tgz
```

The USB device is mounted successfully in the target environment, and files can be copied and managed as needed.

Unmount a USB device

Use this task to safely disconnect a USB device from a sysadmin or XR virtual machine by issuing the appropriate command.

Procedure

Step 1 If you are working in sysadmin-vm, enter the command **usb device operation unmount** to unmount the USB device.

Example:

```
sysadmin-vm:0_RP0#usb device operation unmount  
Fri Jul 13 09:25:24.531 UTC success usb unmounted
```

Step 2 If you are working in XR-vm, enter the command **unmount disk2:** to unmount the USB device.

Example:

```
RP/0/RP0/CPU0:ios#unmount disk2:  
Fri Jul 13 14:56:46.393 IST  
disk2: unmounted successfully.
```

The USB device is safely disconnected from the selected virtual machine and ready for physical removal.



CHAPTER 12

Fault Profiles

- [Fault profiles, on page 169](#)
- [Configure a fault profile, on page 170](#)

Fault profiles

A fault profile is a configuration construct that

- groups fault types generated in a system,
- assigns user-defined severity to each fault, and
- specifies actions such as creating, modifying, or deleting fault profiles and associated alarms

In a system, the default fault list captures all possible faults and their default severity values. The default severity applies when no custom fault profile is attached. Fault profiles can be created for specific scopes, such as the system, line card, node, or port, to override default severity and actions.

If a system profile is attached but you need a different fault profile for a node, you can create a node profile and attach it. The node then inherits its attached node profile properties.

Available severity levels:

- Major
- Minor
- Critical
- Non Faulted
- Non Reported

The defined set of actions for a fault profile are:

- Create and delete a fault profile
- Add alarms to a fault profile
- Remove alarms from a fault profile
- Modify severity of alarm in an existing profile

Best practice for configuring fault profiles

Follow these best practices when configuring fault profiles on Cisco NCS 1004.

- Configure fault profiles only for data path alarms such as Optics, Coherent DSP, Ethernet, and ODU alarms.
- Avoid attempting fault profiling at the port level, as it is not supported.
- Limit the number of fault profiles to a maximum of 61 to prevent configuration errors.

Configure a fault profile

Create a fault profile and apply it to specific hardware or software subsystems so that alarms are triggered with the desired severity and behavior.

Fault profiles allow you to tailor fault management for individual systems, nodes, or propagation domains within Cisco IOS XR devices. By customizing alarm tags and severity levels, you can control how the system responds to faults in various components.

Follow these steps to configure and apply a fault profile:

Procedure

-
- Step 1** Enter the configuration mode.
- Step 2** Define a fault profile with a desired name and attributes using the command **fault-profile *fault_name* fault identifier subsystem XR fault-type { ethernet | sdh_controller | sonet | HW_OPTICS | G709 | CPRI | OTS } fault-tag *alarm_name* severity { sas | nsas } severity_level**. and commit the configuration.
- Step 3** Apply the fault profile at system or node level using the command **fault-profile *fault-name* apply rack *rack_id* slot { ALL | LC }**.
- Step 4** Commit the application and exit the configuration mode.
-

The fault profile is configured and applied, and alarms are generated according to the severity and locations specified.



APPENDIX **A**

Supported SNMP MIBs on NCS 1004

Supported SNMP MIBs

NCS 1004 supports the following SNMP MIBs:

- CISCO-AM-SNMP-MIB
- CISCO-CONFIG-MAN-MIB
- CISCO-FLASH-MIB
- CISCO-ENTITY-REDUNDANCY-MIB
- CISCO-SYSTEM-MIB
- CISCO-ENTITY-ASSET-MIB
- EVENT-MIB
- DISMAN-EXPRESSION-MIB
- CISCO-FTP-CLIENT-MIB
- NOTIFICATION-LOG-MIB
- CISCO-RF-MIB
- RADIUS-AUTH-CLIENT-MIB
- RADIUS-ACC-CLIENT-MIB
- IEEE8023-LAG-MIB
- CISCO-TCP-MIB
- UDP-MIB
- CISCO-BULK-FILE-MIB
- CISCO-CONTEXT-MAPPING-MIB
- CISCO-OTN-IF-MIB
- CISCO-ENHANCED-MEMPOOL-MIB
- CISCO-PROCESS-MIB

- CISCO-SYSLOG-MIB
- ENTITY-MIB
- CISCO-ENTITY-FRU-CONTROL-MIB
- CISCO-IF-EXTENSION-MIB
- RMON-MIB
- HC-RMON-MIB
- CISCO-OPTICAL-MIB
- CISCO-ENTITY-SENSOR-MIB
- LLDP-MIB

The following table provides more information about SNMP MIBs and related documentation links.

Task	Link
Determine the MIB definitions	SNMP Object Navigator
Configure SNMP	Configure SNMP
Understand SNMP best practices about the recommended order of SNMP query, maximum cache hit, and SNMP retry and timeout recommendation	SNMP Best Practices

Make sure that you configure snmp-server community as the SystemOwner to have the admin-plane parameters to appear to entity MIB. The parameters of fans and power supply units are examples of admin-plane parameters.