Configuration Guide for Cisco NCS 1004, IOS XR Release 7.1.1

First Published: 2019-10-11

Last Modified: 2024-04-12

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

CONTENTS

CHAPTER 1 **New and Changed Information** 1

New and Changed Information 1

CHAPTER 2 **Cisco NCS 1004 Overview** 5

Cisco NCS 1004 Features 5

Interoperability with Cisco NCS 1001 5

Supported Line Cards 6

CHAPTER 3 **Configuring the Card Mode** 7

1.2T and 1.2TL Line Cards 7

Card Modes 7

Sub 50G Configuration 7

Supported Data Rates 8

Configuring the Card Mode 10

Regeneration Mode 13

Configuring the Card in Regen Mode 14

Verifying the Regen Mode 14

Configuring the BPS 14

CHAPTER 4 **Configuring Controllers** 17

AINS 17

AINS States 17

Soak Time Period 18

Configuring AINS 18

Disabling AINS 19

Displaying the AINS Configuration 19

FEC 23
FEC States for Ethernet Controller 23
Configuring FEC on the Ethernet Controller 25
Laser Squelching 26
Idle Insertion 27
LLDP Drop 28
Configuring LLDP Drop 29
Verifying the Status of LLDP Drop 30
Link Layer Discovery Protocol (LLDP) Support on Management Interface 31
Loopback 35
Restore Factory Settings 37
Headless Mode 39
Trail Trace Identifier 39
Chromatic Dispersion 41
Frequency 42
Pseudo Random Binary Sequence 43
Configuring Pseudo Random Binary Sequence 43
Verifying PRBS 44
Viewing PRBS Performance Monitoring Parameters 44

CHAPTER 5 **Performance Monitoring** 47
Configuring PM Parameters 47

CHAPTER 6 **IP Access Lists** 57
IP Access List 57
Configuring an IP Access List 58
Verifying ACLs 59

CHAPTER 7 **Layer 1 Encryption** 61
IKEv2 Overview 62
OTNSec Encryption Overview 64
Prerequisites 65
Limitations 65
Configuration Workflow 65

Configuring an IKEv2 Proposal 67
Configuring an IKEv2 Policy 68
Configuring a Keyring 69
Configuring a IKEv2 Profile 70
Configuring an OTNSec Policy 70
Configuring the GCC Interface 71
Configuring OTNSec on ODU4 Controllers 72
Configuration Example 73
Verification 76
Troubleshooting 76
You May Be Interested In 77

CHAPTER 8 GMPLS UNI for Packet and Optical Integration 79
Understanding GMPLS UNI 79
Use Case Overview 80
Prerequisites 80
Limitations 81
Configuration Workflow 81
Configure LMP and Alien Wavelength on NCS 2000 Node Using CTC 81
Configure LMP and Alien Wavelength on NCS 2000 Node Using CTC for Signaled Unnumbered Circuit 84
Retrieve Ifindex from NCS 2000 Node 86
Configure LMP on Cisco NCS 1004 Node 86
Configure RSVP on NCS 1004 Node 88
Configure MPLS Tunnel on a NCS 1004 Node for Numbered Circuit 89
Verification 90
General Troubleshooting 97
You May Be Also Interested In 97

CHAPTER 9 Understanding Remote Node Management Using GCC 99
Limitations 100
Supported Protocols 100
Enable the GCC Interface 101
Configure the GCC Interface 101

Configure Static Routes Over the GCC Interface 102
Configure OSPF Routes Over the GCC Interface 103
iBGP Support Using GCC 104
Restrictions for iBGP Support Using GCC 104
Enabling the GCC Interface 104
Configuring the Management Interface 105
Configuring the Loopback Interface 105
Configuring the GCC interface 105
Verifying iBGP Support Using GCC 106
Use Case - iBGP Support Using GCC Configuration 107

CHAPTER 10 **Smart Licensing** 111
Understanding Smart Licensing 111
Benefits of Smart Licensing 114
Licensing in NCS 1004 114
Software Entitlements of Cisco NCS 1004 114
Configure Smart Licensing 115
Creating a Token 116
Verifying Smart Licensing Configuration 117
License Registration 120

CHAPTER 11 **USB Device Automount** 123
USB Automount 123
Mount USB Device 123
Unmount USB Device 124

CHAPTER 12 **Fault Profiles** 125
Fault Profiles 125
Limitations of Fault Profiles 126
Tasks for Configuring Fault Profiles 126
Configure Fault Profiles 126

CHAPTER 13 **Implementing Host Services and Applications** 129
HTTP Client Application 129

APPENDIX A SNMP 131

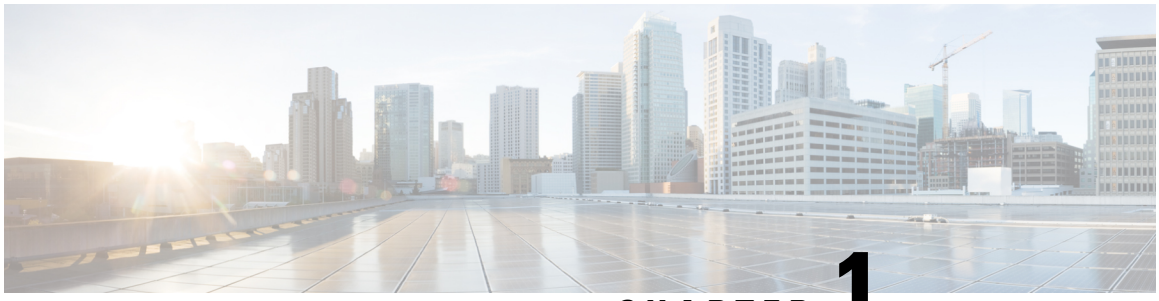

CHAPTER 1

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 configuration guide for Release 7.1.1, and lists where the features are documented.

Table 1: New and Changed Features - R7.1.1

Feature	Description	Where Documented
Fault Profiles	The Fault Profiling feature enables the user to create a unique fault profile for faults on the system or the line card. Each fault profile can contain one or more faults with user-defined severities.	Fault Profiles
Air Filter	The air filter removes dust from the air drawn into the chassisby the fan units. If the air filters are damaged,dirty, or clogged with dust, they must be replaced with a new air filter. Cisco NCS 1004 air filter has the following components: • Two air filter side brackets • One air filter frame • One air filter	Air Filter

Feature	Description	Where Documented
Sub 50G configuration	The sub 50G or coupled mode can be configured on the 1.2T card only in the muxponder mode. The supported trunk data rates are 150G, 250G, 350G, 450G, and 550G.	Sub 50G Configuration
Remote Node Management Using GCC	The remote node management feature allows the user to remotely manage the NCS 1004 nodes over the General Communication Channel (GCC) interface. The remote nodes that are not connected to the management network over the Ethernet interface can be managed over the GCC interface. This feature supports remote management of up to eight nodes in hub topology and up to two nodes in linear topology. GCC2 and GCC0 are supported in NCS 1004.	Remote Node Management Using GCC
Regeneration Mode	The 1.2T card can be configured in Regeneration (regen) mode. In regen mode, only trunk optics Controller and CoherentDSP controllers are created.	Regeneration Mode
PRBS	Pseudo Random Binary Sequence (PRBS) feature enables you to perform data integrity checks between the NCS1004 trunk links without enabling the actual client traffic.	PRBS
Ethernet Statistics	The ethernet statistics are displayed for the current and historical performance monitoring parameters of ethernet controller in 30 second intervals.	Configuring PM Parameters

Feature	Description	Where Documented
LLDP on Management Port	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 can be used to learn about the neighbors and in turn the topology of the devices for Operations, Administration, and Maintenance (OAM) purposes.	Link Layer Discovery Protocol (LLDP) Support on Management Interface
Open Config Enhancements	**CLI Over NETCONF** A new yang model, Cisco-IOS-XR-cli-cfg.yang is defined, which consists of a leaf node called 'cli'. The leaf node can be used to either send or receive the CLI configurations. **OpenConfig Terminal Device** This is a terminal optics device model for managing the terminal systems (client and line side) in a DWDM transport network. **Ethernet Statistics Addition** The user can configure and view the performance monitoring parameters for the Optics, Ethernet, and coherent DSP controllers.	Terminal Device Model

New and Changed Information

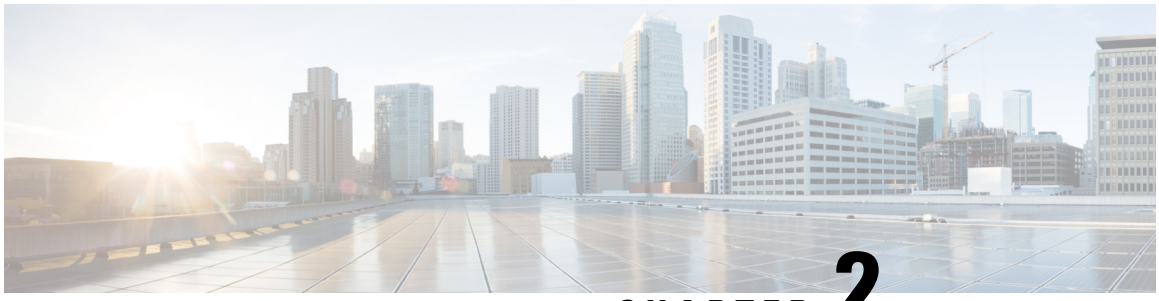

CHAPTER 2

Cisco NCS 1004 Overview

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

- Cisco NCS 1004 Features, on page 5
- Interoperability with Cisco NCS 1001, on page 5
- Supported Line Cards, on page 6

Cisco NCS 1004 Features

Cisco NCS 1004 is a two RU unit that supports up to 4.8 Tbps traffic. The NCS 1004 has two redundant, field replaceable AC and DC power supply units and three redundant, field replaceable fans. It also provides a field replaceable controller card. The NCS 1004 has SSD disks both on board the chassis and on the controller card for resiliency. Each NCS 1004 chassis provides four line cardslots. Each NCS 1004 slot can host a line card. See Supported Line Cards, on page 6 for more information.

The NCS 1004 delivers the following benefits:

- Transport of any trunk rate from 150 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 that is optimized for performance for metro, long-haul, and submarine applications.
- Support for up to 350,000 ps/nm of residual chromatic dispersion compensation.
- Transport of 100GE and OTU4 client rates on the same platform through software provisioning.
- 600G DWDM, which provides unparalleled scale and density. 64 channels of 600G at 75 GHz providing 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 with Cisco NCS 1001

When the Cisco NCS 1001 with Protection Switching Module (PSM) configured as non-revertive, interoperates with Cisco NCS 1004, traffic loss may occur. After the traffic has switched from the working to the protect

path, do not perform a manual switch for 120 seconds. If you perform a manual switch, and the protect path fails, traffic loss of up to 13 seconds can occur.

Supported Line Cards

The following 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 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 format, baud-rate, and forward error correction. The trunk ports are software configurable. The line card supports 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 format can be QPSK, 8 QAM, 16 QAM, 32 QAM, or 64 QAM.
- Hybrid modulations formats can be configured through 1/128 bits/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.

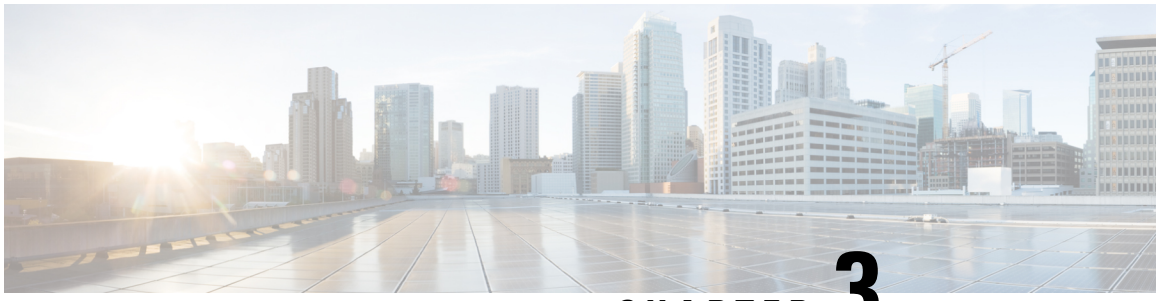

CHAPTER 3

Configuring the Card Mode

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

- 1.2T and 1.2TL Line Cards, on page 7

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

The line cards support module and slice configurations.

The line cards have two trunk ports (0 and 1) and 12 client ports (2 through 13) each. You can configure the line card in two modes:

- Muxponder—In this mode, both trunk ports are configured with the same trunk rate. The client-to-trunk mapping is in a sequence.
- Muxponder slice—In this mode, each trunk port is configured independent of the other with different trunk rates. The client-to-trunk mapping is fixed. For Trunk 0, the client ports are 2 through 7. For Trunk 1, the client ports are 8 through 13.

Sub 50G Configuration

You can configure the sub 50G or coupled mode on the line card only in the muxponder mode. The following table displays the port configuration for the 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)
150G	300G	1.2T, 1.2TL	0, 1	2	3	4
250G	500G	1.2T	0, 1	2, 3	4	5, 6

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)
350G	700G	1.2T	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

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

Coupled Mode Restrictions

The following restrictions apply to the coupled mode configuration:

- Both trunk ports must be configured with the same bits-per-symbol or baud rate and must be sent over same fiber and direction.
- The chromatic dispersion must be configured to the same value for both trunk ports.
- When trunk internal loopback is configured, it must be done for both trunk ports. Configuring internal loopback on only one trunk results in traffic loss.
- 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 R7.1.1 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. 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
200	1.2T	100GE, OTU4	0, 1	2, 3, 4, 5
300	1.2T	100GE, OTU4	0, 1	2, 3, 4, 5, 6, 7
400	1.2T	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

Trunk Data Rate	Card Support	Client Data Rate (100GE, OTU4)	Trunk Ports	Client Ports
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	100, OTU4	0	2
200	1.2T	100, OTU4	0	2, 3
300	1.2T	100, OTU4	0	2, 3, 4
400	1.2T	100, OTU4	0	2, 3, 4, 5
500	1.2T	100, OTU4	0	2, 3, 4, 5, 6
600	1.2T	100, 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	100, OTU4	1	8
200	1.2T	100, OTU4	1	8, 9
300	1.2T	100, OTU4	1	8, 9, 10
400	1.2T	100, OTU4	1	8, 9, 10, 11
500	1.2T	100, OTU4	1	8, 9, 10, 11, 12
600	1.2T	100, 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
200G	27%	2	4.40625	31.51	69.43
300G	27%	2.8984375	6	34.7175497	71.8681352
400G	27%	3.8671875	6	46.2900663	71.8197392
500G	27%	4.8281250	6	57.8625828	71.9068991
600G	15%	5.2578125	-	-	71.9552971

To configure the BPS, see Configuring the BPS, on page 14.

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** | **100G150G** | **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**

hw-module location *location* **mxponder-slice trunk-rate** { **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
```

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

HundredGigECtrlr0/2/0/2  ODU40/2/0/0/1               100                   0
HundredGigECtrlr0/2/0/3  ODU40/2/0/0/2               100                   0
HundredGigECtrlr0/2/0/4  ODU40/2/0/0/3               100                   0
HundredGigECtrlr0/2/0/5  ODU40/2/0/0/4               100                   0
HundredGigECtrlr0/2/0/6  ODU40/2/0/0/5               100                   0
HundredGigECtrlr0/2/0/7  ODU40/2/0/1/1                 0                 100
HundredGigECtrlr0/2/0/8  ODU40/2/0/1/2                 0                 100
HundredGigECtrlr0/2/0/9  ODU40/2/0/1/3                 0                 100
HundredGigECtrlr0/2/0/10 ODU40/2/0/1/4                 0                 100
HundredGigECtrlr0/2/0/11 ODU40/2/0/1/5                 0                 100
```

The following is a sample ouput 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

HundredGigECtrlr0/1/0/2    ODU40/1/0/0/1             100                   0
HundredGigECtrlr0/1/0/3    ODU40/1/0/0/2             100                   0
HundredGigECtrlr0/1/0/4    ODU40/1/0/0/3             100                   0
HundredGigECtrlr0/1/0/5    ODU40/1/0/0/4             100                   0
HundredGigECtrlr0/1/0/6    ODU40/1/0/0/5             100                   0
HundredGigECtrlr0/1/0/7    ODU40/1/0/0/6              50                  50
HundredGigECtrlr0/1/0/8    ODU40/1/0/1/1               0                 100
HundredGigECtrlr0/1/0/9    ODU40/1/0/1/2               0                 100
HundredGigECtrlr0/1/0/10   ODU40/1/0/1/3               0                 100
HundredGigECtrlr0/1/0/11   ODU40/1/0/1/4               0                 100
HundredGigECtrlr0/1/0/12   ODU40/1/0/1/5               0                 100
```

The following is a sample ouput 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 ouput 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 ouput 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
```

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
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                                     : 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
```

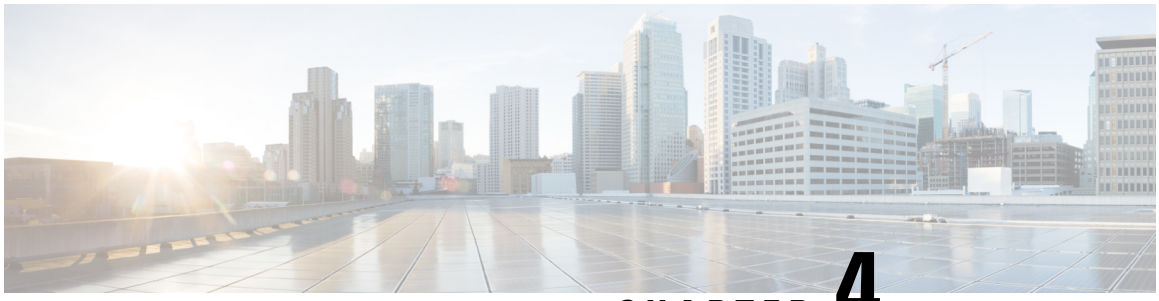

CHAPTER 4

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.

- AINS, on page 17
- FEC, on page 23
- Laser Squelching, on page 26
- Idle Insertion, on page 27
- LLDP Drop, on page 28
- Link Layer Discovery Protocol (LLDP) Support on Management Interface, on page 31
- Loopback, on page 35
- Restore Factory Settings, on page 37
- Headless Mode, on page 39
- Trail Trace Identifier, on page 39
- Chromatic Dispersion, on page 41
- Frequency, on page 42
- Pseudo Random Binary Sequence, on page 43

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.

State	Description
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 runring 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
```

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.
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 HundredGigECtrlr 0/0/0/4
Thu Aug  8 15:41:20.857 IST
Operational data for interface HundredGigECtrlr0/0/0/4:

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
```

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
```

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
```

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 2:

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**

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 HundredGigECtrlr 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 HundredGigECtrlr 0/1/0/10
Fri Feb 22 19:50:08.328 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: 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
```

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
ÿûÿûÿûÿýbootlogd: ioctl(/dev/pts/2, TIOCCONS): Device or resource busy
/sbin/restorecon:  lstat(/etc/adjtime) failed:  No such file or directory
Configuring network interfaces... done.
Starting system message bus: dbus.
Starting OpenBSD Secure Shell server: sshd
sshd start/running, process 1739
Starting rpcbind daemon...done.
Starting random number generator daemonUnable to open file: /dev/tpm0
.
Starting system log daemon...0
Starting kernel log daemon...0
tftpd-hpa disabled in /etc/default/tftpd-hpa
Starting internet superserver: xinetd.
net.ipv4.ip_forward = 1
Libvirt not initialized for container instance
```

```
Starting crond: OK
SIOCADDRT: File exists

DBG_MSG: platform type is 0
[*] ima_policy have loaded, or IMA policy file does not exist
Start serial incoming on , Clearing ..
RP/0/RP0/CPU0:Sep 24 09:38:44.284 UTC: fpd-serv[256]: %PKT_INFRA-FM-3-FAULT_MAJOR :
ALARM_MAJOR :FPD-NEED-UPGRADE :DECLARE :0/PM0:

This (D)RP Node is not ready or active for login /configuration
.......
........
........

ios con0/RP0/CPU0 is now available

Press RETURN to get started.

!!!!!!!!!!!!!!!!!! NO root-system username is configured. Need to configure root-system
username. !!!!!!!!!!!!!!!!!!
```

Headless Mode

During process restarts, CPU reload, or removal of CPU, the NCS 1004 operates in headless mode for up to 72 hours. During this time, traffic is not impacted, although the control plane is not up and running. Fault propagation continues to operate for failures on client and trunk ports. However, you cannot provision anything nor view operational data with a non-functional CPU. Performance monitoring data based on 15 minutes and 24 hour intervals is not supported with a non-functional CPU.

Trail Trace Identifier

The Trail trace identifier (TTI) feature helps you to identify the signal from the source to the destination within the network. You can configure the TTI sent or expected string only in ASCII string format. When the expected TTI string does not match the received TTI trace string, the controller goes down and the OTUK-TIM alarm is raised. To configure TTI on the coherent DSP controllers, use the following commands:

configure

controller coherentDSP *R/S/I/P* **tti** {**sent | expected**} **ascii** *tti-string*

commit

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
                                              : 00000000000000000000000000000000
OTU TTI Received
     OPERATOR SPECIFIC  ASCII                 : 1234
                                              :
     OPERATOR SPECIFIC HEX                    : 31323334000000000000000000000000
                                              : 00000000000000000000000000000000
OTU TTI Expected
     OPERATOR SPECIFIC  ASCII                 : 5678
                                              :
     OPERATOR SPECIFIC HEX                    : 35363738000000000000000000000000
                                              : 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	-10,000 to 100,000 ps/nm
	3 < BPS <= 4	1.2T	-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

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.

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/I/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 3: 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.

PM Parameter	Description
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
```

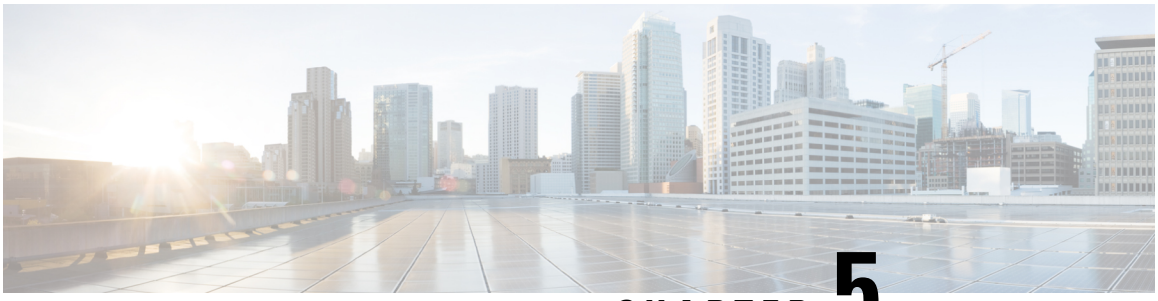

CHAPTER 5

Performance Monitoring

Performance monitoring (PM) parameters are used by service providers to gather, store, set thresholds for, and report performance data for early detection of network issues. You can configure and retrieve PM counters for the various controllers in 30-second, 15-minute, or 24-hour intervals. These parameters simplify troubleshooting operations and enhance data that can be collected directly from the equipment.

- Configuring PM Parameters, on page 47

Configuring PM Parameters

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

To configure PM parameters, use the following commands.

configure

controller *controllertype R/S/I/P* { **pm** { **15-min** | **30-sec** | **24-hour** } { **optics** | **ether** | **pcs** | **fec** | **otn**} { **report** | **threshold** } *value* }

commit

Examples

The following 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
```

The following 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
```

The following is a sample in which performance monitoring parameters of a Coherent DSP controller are configured 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
```

Viewing PM Parameters

To view the performance monitoring parameters for Optics, Ethernet, and Coherent DSP controllers, use this command:

show controllers *controllertype R/S/I/P* **{ pm { current | history } { 30 sec | 15-min | 24-hour } { optics | ether | fec | otn | prbs}** *linenumber* }

Example 1: 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(min)   Threshold(min) (min) Threshold(max)
   Threshold(max) (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
```

Example 2: 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(min)   Threshold(min) (min) Threshold(max)
   Threshold(max) (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
```

Example 3: 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.

Example 4: 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   TCA     Threshold   TCA
                                              (min)    (enable)   (max)    (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
```

Example 5: 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
```

Example 6: Displays the current PCS performance monitoring parameters of the Coherent DSP controller 30-second intervals.

```
RP/0/RP0/CPU0:ios#show controllers hundredGigECtrlr 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#
```

Example 7: 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[06] : 0
BIP[07] : 0
BIP[08] : 0
BIP[09] : 0
BIP[10] : 0
BIP[11] : 0
BIP[12] : 0
BIP[13] : 0
BIP[14] : 0
BIP[15] : 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
FRM-ERR[03] : 0
FRM-ERR[04] : 0
FRM-ERR[05] : 0
FRM-ERR[06] : 0
FRM-ERR[07] : 0
FRM-ERR[08] : 0
FRM-ERR[09] : 0
FRM-ERR[10] : 0
FRM-ERR[11] : 0
FRM-ERR[12] : 0
FRM-ERR[13] : 0
FRM-ERR[14] : 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[03] : 0
BAD-SH[04] : 0
BAD-SH[05] : 0
BAD-SH[06] : 0
BAD-SH[07] : 0
BAD-SH[08] : 0
BAD-SH[09] : 0
BAD-SH[10] : 0
BAD-SH[11] : 0
BAD-SH[12] : 0
BAD-SH[13] : 0
BAD-SH[14] : 0
BAD-SH[15] : 0
BAD-SH[16] : 0
BAD-SH[17] : 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#
```

Clearing PM Parameters

To clear the performance monitoring parameters for Ethernet and Coherent DSP controllers, use this command:

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

Example 1: 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   TCA     Threshold   TCA
                                              (min)    (enable)   (max)    (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   TCA      Threshold   TCA
                                                (min)     (enable)   (max)     (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
```

Example 2: Clears the PM parameters on the Ethernet controller.

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

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 encap    = 0
    Input drop other            = 0

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

    Input MIB giant             = 0
```

```
    Input MIB jabber            = 0
    Input MIB CRC               = 0

Egress:
    Output total bytes          = 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.

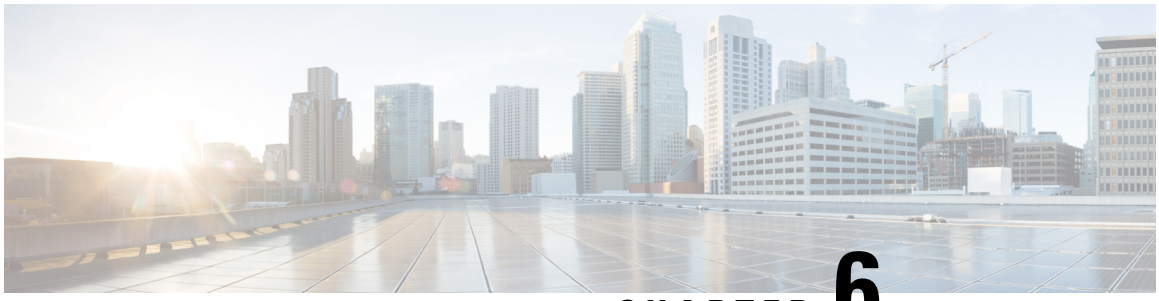

CHAPTER 6

IP Access Lists

This chapter describes how to configure IPv4 and IPv6 Access Control Lists (ACL).

- IP Access List, on page 57

IP Access List

How an IP Access List Works

An access list is a sequential list consisting of permit and deny statements that apply to IP addresses and possibly the upper-layer IP protocols. ACLs are used to permit or deny the flow of packets based on matching criteria of access list parameters and information contained in packets. For it to be in effect, an access list must be created and applied to an interface.

An access list can control traffic arriving or leaving the system, but not traffic originating at the system.

IP Access List Process and Rules

There are two paths for interface packet filtering for ACL configuration:

- Hardware programming path: Hardware programming path is the fast path ACL configuration. The fast path ACL configuration requires Ternary Content Addressable Memory (TCAM) through packet filter Execution Agent.
- Software programming path: Software programming path is the slow path ACL configuration. The slow path ACL configuration requires adding caps to Interface Manager and NetIO.

Use the following process and rules when configuring an IP access list:

- The software tests the source or destination address or the protocol of each packet being filtered against the conditions in the access list, one condition (permit or deny statement) at a time.
- If a packet does not match a statement in the access list, it is tested against the next statement in the list.
- If a packet matches an access list statement, the remaining statements in the list are skipped, and the packet is permitted or denied as specified in the matched statement.
- If the access list denies the address or the protocol, the software discards the packet and returns an Internet Control Message Protocol (ICMP) Host Unreachable message. ICMP is configurable in the Cisco IOS XR software.
- If no conditions match, the software drops the packet because each access list ends with an unwritten or implicit deny statement.

- The access list should contain at least one permit statement; otherwise, all packets are denied.
- The software stops testing the conditions after the first match; so, the order of the conditions is critical. The same permit or deny statements specified in a different order could result in a packet being passed under one circumstance and denied in another circumstance.
- Only one access list per interface, per protocol, per direction is allowed.
- Inbound access lists process packets arriving at the system. Incoming packets are processed before being routed to an outbound interface. An inbound access list is efficient as it saves the overhead of routing lookups if the packet is to be discarded because it is denied by the filtering tests. If the packet is permitted by the tests, it is then processed for routing. For inbound lists, "permit" means continue to process the packet after receiving it on an inbound interface; "deny" means discard the packet.
- Outbound access lists process packets before they leave the system. Incoming packets are routed to the outbound interface and then processed through the outbound access list. For outbound lists, "permit" means send it to the output buffer; "deny" means discard the packet.
- An access list cannot be removed if that access list is being applied by an access group in use. To remove an access list, remove the access group that is referencing the access list and then remove the access list.
- An access-list must be created first before it can actually be applied on the management interface using access-group command.
- ACLs apply only on management interfaces and not on any other type of interfaces or controllers.

Statistics collections are also divided into fast path packets and slow path packets. ACLs information is stored as a global data on the route processor.

Support of IP Access list in NCS 1004:

NCS 1004 supports the following:

- Ingress ACL for both IPv4 and IPv6.
- Slow packet path for Management Interface.
- Egress ACL: Self-Originated Packet is not supported by ACL, because this is already controlled by the user. Only forwarded packets or traffic classify under ACL. This rule is applicable for both IPv4 and IPv6 ACL.

Configuring an IP Access List

To configure the ACL, use the following commands at the IPv4 or IPv6 interface:

configure

interface *interface-type Rack/Slot/Instance/Port*

ipv4 | ipv6 access-group *access-list-name* **{ingress | egress}**

commit

Example

```
interface MgmtEth0/RP0/CPU0/0
ipv4 address 10.1.1.1 255.255.255.0
ipv6 address 1000::1/64
ipv4 access-group IPV4_ICMP_DENY ingress
```

```
ipv4 access-group IPV4_ROUTER_FWD_TELNET_TRAFFIC_DENY egress
ipv6 access-group IPV6_SSH_DENY ingress
ipv6 access-group IPV6_ROUTER_FWD_TELNET_TRAFFIC_DENY egress
```

Sample Configuration for IPv4 Access Lists

```
ipv4 access-list IPV4_ICMP_DENY
10 deny icmp any any
20 permit ipv4 any any
!
ipv4 access-list IPV4_ROUTER_FWD_TELNET_TRAFFIC_DENY
10 deny tcp any any eq telnet
20 permit ipv4 any any
!
```

Sample Configuration for IPv6 Access Lists

```
ipv6 access-list IPV6_SSH_DENY
10 deny tcp any any eq ssh
20 permit ipv6 any any
!
ipv6 access-list IPV6_ROUTER_FWD_TELNET_TRAFFIC_DENY
10 deny tcp any any eq telnet
20 permit ipv6 any any
!
```

Verifying ACLs

The following examples verify the number of packets filtered by the respective ACLs:

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)
```

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)
```

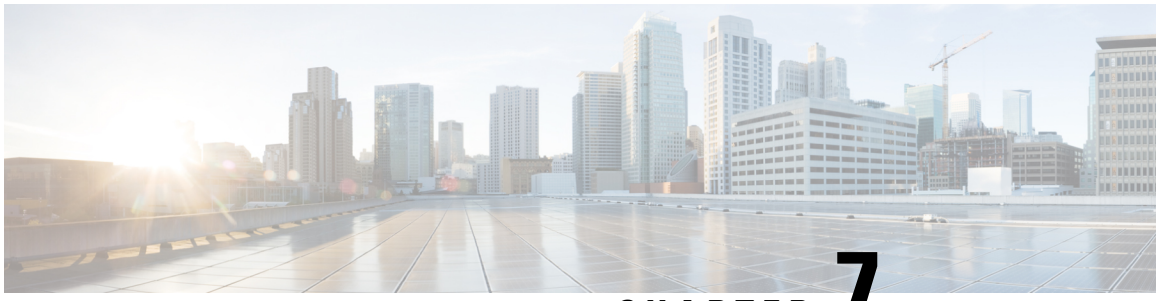

CHAPTER 7

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 4: 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) 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 1: OTNSec Site-to-Site Example and Components

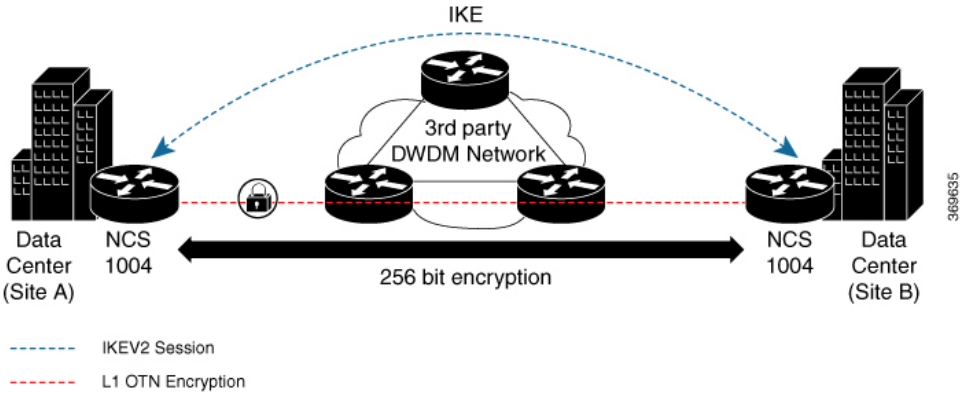

The recommended deployment is to have a single IKEv2 session running over a GCC2 channel per trunk port which creates the child SAs for each of the OTNSec controllers that are configured on the trunk port.

- IKEv2 Overview, on page 62
- OTNSec Encryption Overview, on page 64
- Prerequisites, on page 65
- Limitations, on page 65
- Configuration Workflow, on page 65
- Configuration Example, on page 73
- Verification, on page 76
- Troubleshooting, on page 76
- You May Be Interested In, on page 77

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 2: IKEv2 Exchanges

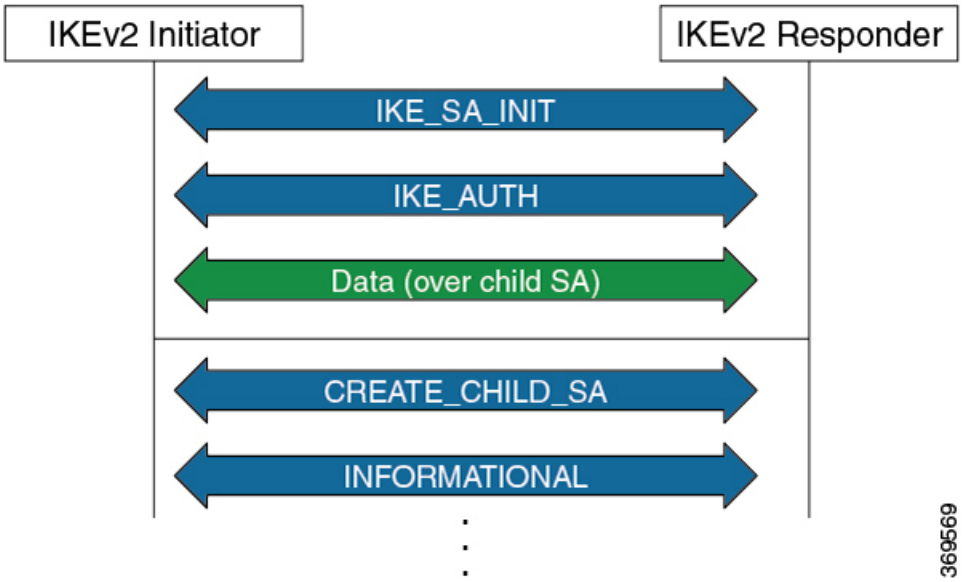

IKEv2 is defined in RFC 7296 and consists of the following constructs:

- **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.

- **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.

- **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**

      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 3: L1 Encryption Workflow

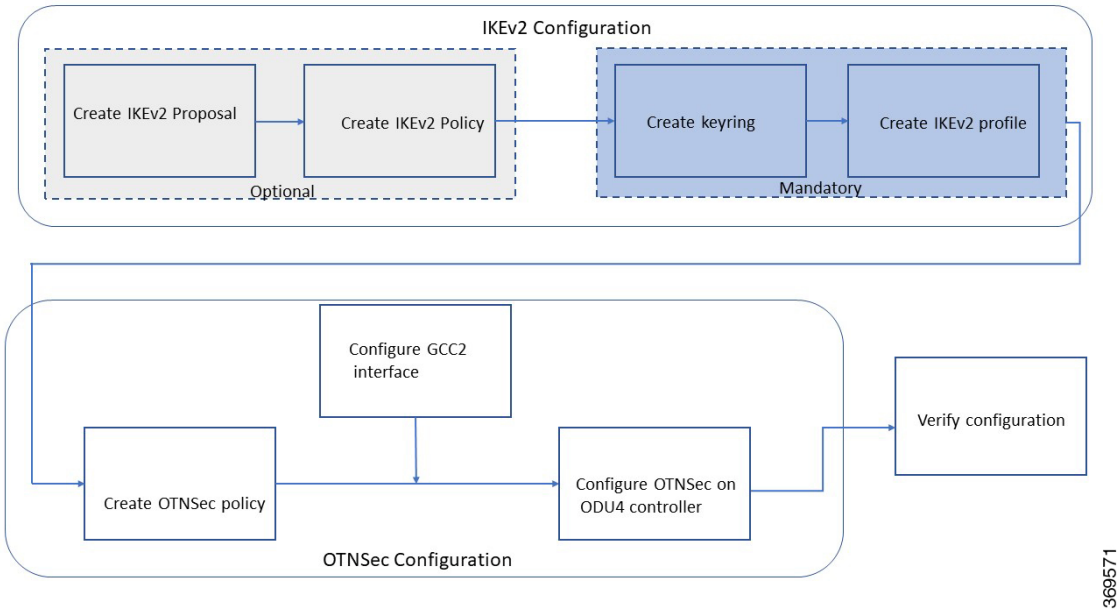

Table 5: Workflow for Configuring IKEv2 and OTNSec Encryption on NCS 1004

Workflow Sequence	Details
IKE Configuration	

Configuring an IKEv2 Proposal, on page 67	(Optional) Configure an IKEv2 proposal manually;otherwise, the default IKEv2 proposal is used in the default IKEv2 policy.
	The default IKEv2 proposal requires no configuration and is a collection of commonly used transforms types, which are as follows:
	`encryption cbc-aes-256`
	`integrity sha512, sha384`
	`prf sha512, sha384`
	`dh 19, 20, 21`

Configuring an IKEv2 Policy, on page 68	(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 69	Configure a keyring as the local or remote authentication method is a preshared key.

Workflow Sequence	Details
Configuring a IKEv2 Profile, on page 70	Configure an IKEv2 profile. **Note** • 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 70	(Optional) Configure the OTNSec policy.
Configuring the GCC Interface, on page 71	Configure the GCC2 interface.
Configuring OTNSec on ODU4 Controllers, on page 72	Configure the ODU4 controller that is mapped to the HundredGigE controller .
Verification	
Verification, on page 76	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***} {password***encrypted 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
   Local PSK                          : Configured
   Remote PSK                         : Configured
```

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 120
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 ODUC4 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 4: Configuration Schema

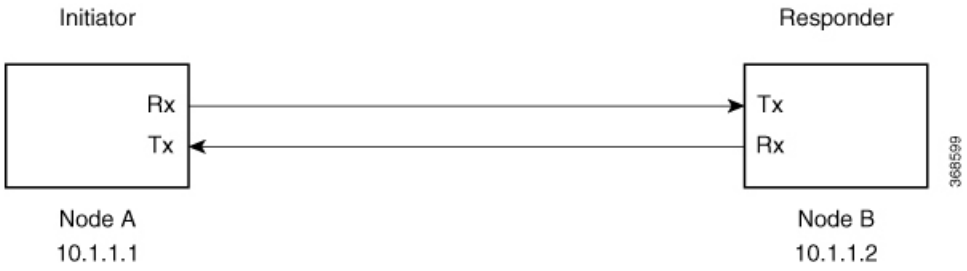

The configuration on Node A is displayed below.

Node A (Initiator)
Keyring
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
IKEv2 profile
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 600 RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#commit RP/0/RP0/CPU0:SITE-A(config-ikev2-profile-IP1)#exit
OTNSec policy

Node A (Initiator)

```
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
```

GCC interface

```
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
```

OTNSec on ODU4 controller

```
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
```

The configuration on Node B is displayed below.

Node B (Responder)

Keyring

```
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
```

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 600
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 ODUC4 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 6: 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.

You May Be Interested In

- For more information about IKEv2, see RFC 7296.
- For more information about NCS 1004, see the NCS 1004 datasheet.

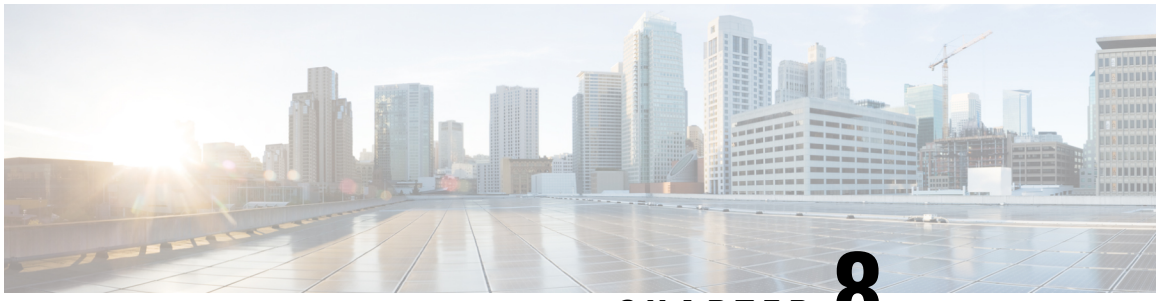

CHAPTER 8

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 79
- Use Case Overview, on page 80
- Prerequisites, on page 80
- Limitations, on page 81
- Configuration Workflow, on page 81
- Verification, on page 90
- General Troubleshooting, on page 97
- You May Be Also Interested In, on page 97

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 Configure LMP on Cisco NCS 1004 Node 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 5: GMPLS UNI Reference Model

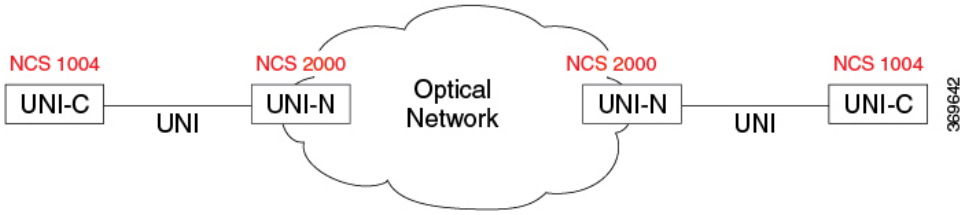

UNI-C is the client or packet or router node; for example, NCS 1004 nodes. UNI-N is the network or optical node; for example, NCS 2000 nodes.

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.
   - Configure LMP and Alien Wavelength on NCS 2000 Node Using CTC , on page 81
2. Retrieve Ifindex from NCS 2000 Node, on page 86

Configurations on the node:

1. Configure LMP on Cisco NCS 1002 Node.
   - Configure LMP on Cisco NCS 1004 Node , on page 86
2. Configure RSVP on NCS 1004 Node, on page 88
3. Configure MPLS Tunnel on a NCS 1002 Node.
   - Configure MPLS Tunnel on a NCS 1004 Node for Numbered Circuit, on page 89

Configure LMP and Alien Wavelength on NCS 2000 Node Using CTC

This procedure creates a static LMP link to establish connectivity between a NCS 2000 node and a NCS 1004 node. The LMP creation wizard in CTC provides the capability to select source and destination endpoints of the LMP link, optical parameters, and alien wavelength settings.

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 that 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 newly created LMP link appears in the LMP table in CTC.

Configure LMP and Alien Wavelength on NCS 2000 Node Using CTC for Signaled Unnumbered Circuit

This procedure creates a static LMP unnumbered link to establish connectivity between a NCS 2000 node and a NCS 1004 node. The LMP creation wizard in CTC provides the capability to select source and destination endpoints of the LMP link, optical parameters, and alien wavelength settings.

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.
- 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 In 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.

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 Configure LMP on Cisco NCS 1004 Node 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 Configure LMP on Cisco NCS 1004 Node .

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 86 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 7: 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)

                  State          LMP          GMPLS tunnel
Interface      Admin  Oper    adjacency    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
      Fspec: 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: ncs1k, 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. http://*ip-address-of-head-node*/diagnostics/wson

You May Be Also Interested In

- GMPLS UNI commands: Cisco IOS XR MPLS Command Reference.
- GMPLS Restoration and Reversion

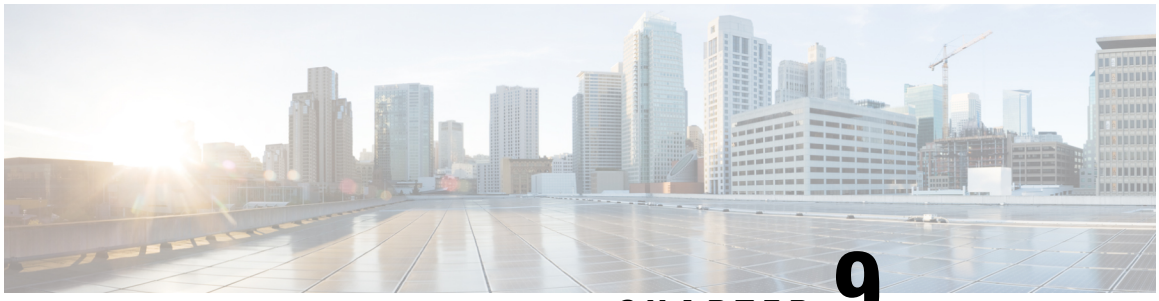

CHAPTER 9

Understanding Remote Node Management Using GCC

The remote node management feature allows you to remotely manage NCS 1004 nodes over the General Communication Channel (GCC) interface. The remote nodes that are not connected to the management network over the Ethernet interface can be managed over the GCC interface. This feature supports remote management of up to eight nodes in hub topology and up to two nodes in linear topology.

Figure 6: Remote Node Management in Linear Topology

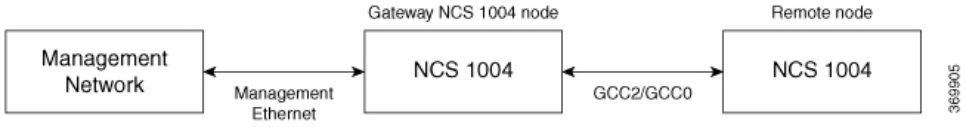

The remote nodes can be dynamically discovered over the GCC interface using OSPF. The connectivity to the management network can be achieved using OSPF and static routes.

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.

- Limitations, on page 100
- Supported Protocols, on page 100
- Enable the GCC Interface, on page 101
- Configure the GCC Interface, on page 101
- Configure Static Routes Over the GCC Interface, on page 102
- Configure OSPF Routes Over the GCC Interface, on page 103
- iBGP Support Using GCC, on page 104

Limitations

- gRPC is not supported over the GCC interface. Therefore, Open Config and streaming telemetry are not supported over the GCC interface.
- Only the Tx and Rx packet count information are available in GCC statistics.
- The devices can be remotely managed over the GCC interface only when they are connected to the management network through GCC. Therefore, initial provisioning and bringing up of the GCC interface must be performed either through the console or management Ethernet interface.
- The following headless or high availability events at the intermediate nodes may affect remote node management of subsequent nodes:
  - Reload of the route processor
  - Reload of IOS XR
  - Restart of the driver process
- IP fragmentation is not supported on GCC interface for the SCP protocol. As a workaround, you can apply any of the following configurations to limit the maximum packet size below the fragmentation limit (1454 bytes):
  - Use the **tcp mss** <*maximum segment size*> command (for example, **tcp mss** *1200*) in the global configuration mode. The maximum segment limit is applied to all interfaces.
  - Use the **ipv4 mtu** <*MTU size*> command in the interface configuration mode. The MTU size is applied only to the specified interface.

Supported Protocols

The following protocols are supported over the GCC interface.

- PING
- SSH
- TELNET
- SCP
- TFTP
- FTP
- SFTP
- HTTP
- HTTPS
- OSPF

Enable the GCC Interface

Enable the GCC Interface on 1.2T Card

To enable the GCC2 interface for the 1.2T line card, use the following commands:

```
configure
controller odu4 R/S/I/P/L
gcc2
commit
exit
```

To enable the GCC0 interface for the 1.2T line card, use the following commands:

```
configure
controller CoherentDSP R/S/I/P
gcc0
commit
exit
```

Enable the GCC Interface on OTN-XP Card

Configure the GCC Interface

Configure the GCC Interface on 1.2T Card

To configure the GCC2 interface using the static IP address for the 1.2T card, use the following commands:

configure

interface gcc2 *R/S/I/P/L*

ipv4 address *ipv4-address*

commit

exit

To configure the GCC0 interface using the static IP address for the 1.2T card, use the following commands:

configure

interface gcc0 *R/S/I/P*

ipv4 address *ipv4-address*

commit

exit

Configure the GCC Interface on OTN-XP Card

Examples

The following sample displays how to configure the GCC2 interface using the static IP address on 1.2T 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
!
```

The following sample displays how to configure 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
RP/0/RP0/CPU0:ios(config)#exit
```

The following 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              Up       default
```

The following sample displays how to configure the GCC0 interface using the static IP address on 1.2T or OTN-XP 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
!
```

The following sample displays how to configure 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

To configure the static routes over the GCC interface, use the following commands:

configure

router static address-family ipv4 unicast 0.0.0.0/0 *default-gateway*

exit

Examples

```
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

To configure OSPF routes over the GCC interface, use the following commands:

configure

router ospf *process-id*

router-id *ip-address*

area *area-id*

interface *type R/S/I/P/L*

exit

Examples

The following is a sample to configure OSPF routes over the GCC interface.

Gateway Node:

```
configure
router ospf 1
router-id 192.0.2.89

 area 0

  interface Loopback0

  !

  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
```

iBGP Support Using GCC

The Internal BGP (iBGP) support over GCC allows external devices to exchange BGP routes through management interfaces of NCS1004 system. The NCS 1004 device advertises local networks through BGP and manages these networks using path learnt through BGP. With the iBGP route information, the NCS 1004 devices establish iBGP sessions over GCC to exchange BGP routes.

You can configure VPN routing and forwarding (VRF) on the GCC management interfaces (port 0 and port 1) of the NCS 1004 device. The VRF enables traffic isolation between the management ports (port 0 and port 1).

The GCC2 and GCC0 interfaces are supported in NCS 1004 for 1.2 T line card.

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.

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
```

Use Case - iBGP Support Using GCC Configuration

Consider two NCS 1004 devices R2 and R3 connected through GCC0 interfaces.

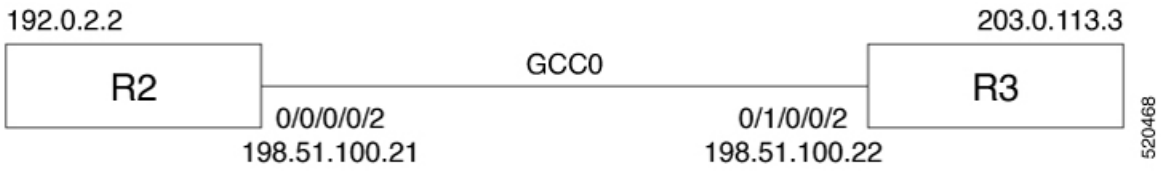

R2 is connected through GCC0 0/0/0/0 interface with IP address of 198.51.100.21 and R3 is connected through GCC0 0/1/0/0 with IP address of 198.51.100.22. The R2 and R3 devices are connected to external devices through management interfaces.

Table 8:

Configuration on R2	Configuration on R3

Global Configuration on R2
 	
`hw-module location 0/0`
` mxponder`
`  trunk-rate 600G`
`  client-rate 100GE`
 	
`vrf transport-vrf`
` address-family ipv4 unicast`
	Global Configuration on R3
	
	`hw-module location 0/0`
	` mxponder`
	`  trunk-rate 600G`
	`  client-rate 100GE`
	
	`vrf transport-vrf`
	` address-family ipv4 unicast`

Interface Configuration on R2
 	
`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`
	Interface Configuration on R3
	
	`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`

Route-policy Configuration on R2
 	
`route-policy PASS-ALL`
`  pass`
`end-policy`
	Router Policy Configuration on R3
	
	`route-policy PASS-ALL`
	`  pass`
	`end-policy`

Static Route Configuration on R2
 	
`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`
	Static Route Configuration on R3
	
	`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`

Configuration on R2	Configuration on R3

BGP Configurationon R2
 	
`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`
`   !`
	BGP Configuration on R3
	
	`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`
	`   !`

BGP Verification on R2
 	
`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`
	BGP Verification on R3
	
	`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`

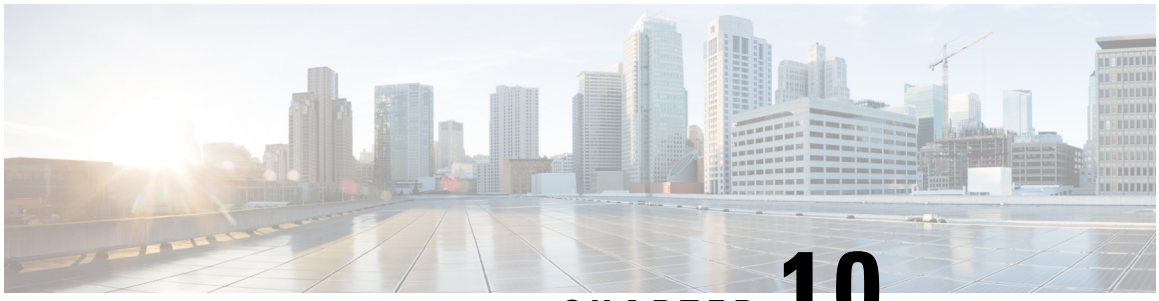

CHAPTER 10

Smart Licensing

This chapter describes the smart licensing configuration on Cisco NCS 1004.

- Understanding Smart Licensing, on page 111
- Configure Smart Licensing, on page 115

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

Cisco Smart Account is an account where all products enabled for Smart Licensing are deposited. Cisco Smart Account allows you to manage and activate your licenses to devices, monitor license use, and track Cisco license purchases. Through transparent access, you have a real-time view into your Smart Licensing products. IT administrators can manage licenses and account users within your organization's Smart Account through the Smart Software Manager.

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

A Virtual Account exists as a sub-account tithing the Smart Account. Virtual Accounts are a customer-defined structure based on organizational layout, business function, geography, or any defined hierarchy. 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 7: Smart Licensing Work Flow

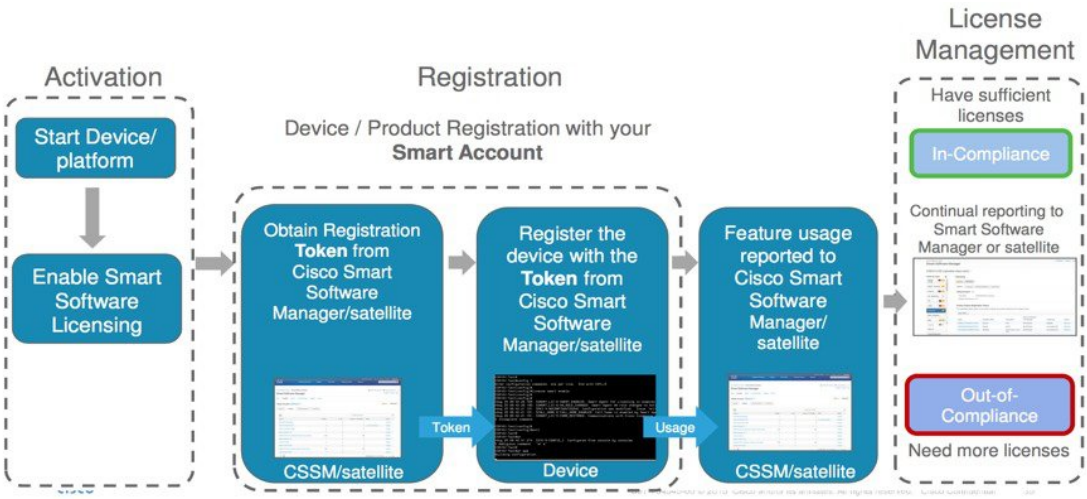

1. **Setting up Smart Licensing**: You can place the order for Smart Licensing, to manage licenses on the Cisco.com portal. You agree to the terms and conditions governing the use and access of Smart Licensing in the Smart Software Manager portal.

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-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 9: Software Entitlements of Cisco NCS 1004

Feature	Software Entitlement
NCS1K4 Smart License-one QSFP28 client	S-NCS1K4-LIC-100G=

Feature	Software Entitlement
NCS1K4 Smart License - one QSFP28 client with encryption	S-NCS1K4-LIC-100X=

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 116.

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 10: 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.

License Authorization Status	Description
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/CPU:ios#license smart register idtoken <idtoken>
```

Step 2 Browse to the URL : https://software.cisco.com/software/csws/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.

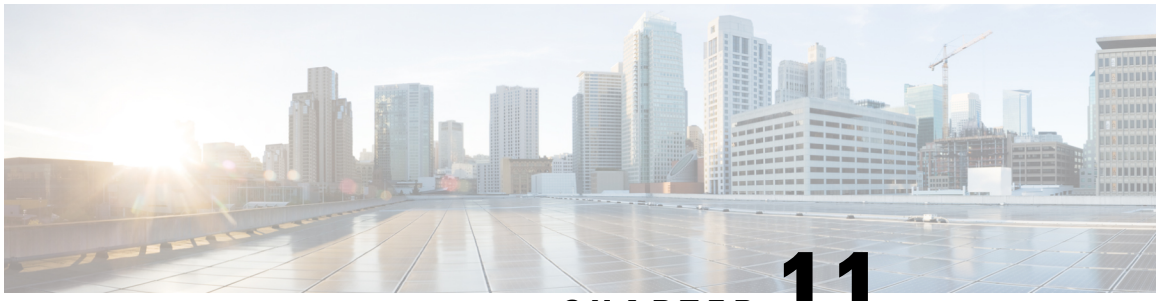

CHAPTER 11

USB Device Automount

This chapter describes the USB automount configuration on NCS 1004.

- USB Automount, on page 123
- Mount USB Device, on page 123
- Unmount USB Device, on page 124

USB Automount

When you insert a USB device in NCS 1004, it is automatically mounted in sysadmin-vm with Read and Write permissions and unmounted in XR by default. The USB device is automatically mounted in sysadmin-vm only when the file system of the USB device is FAT or FAT32. This feature allows you to read from or write files and folders onto the USB device without explicitly mounting it. You can access the mounted USB device as disk2: file system.

Unmount the USB device from sysadmin-vm before you remove it from NCS 1004. Use the **mount** command, if you want to mount the USB device again after you unmount the device but before physically removing it.

You can simultaneously mount the USB device in XR-vm and sysadmin-vm. Unmount the USB device from both sysadmin-vm and XR-vm before you remove it from NCS 1004.

Mount USB Device

To mount the USB device in sysadmin-vm, use the following command:

usb device operation mount

To mount the USB device in XR-vm, use the following command:

unmount disk2: undo

Example

The following is an example of USB mount in sysadmin-vm.

```
sysadmin-vm:0_RP0#usb device operation mount
Fri Jul  13 09:26:00.821 UTC success usb mounted
```

The following is an example of USB mount in XR-vm.

```
RP/0/RP0/CPU0:ios#unmount disk2: undo
Fri Jul 13 14:56:34.326 IST
disk2: mounted successfully.
```

The following is an example of copying the file to the USB device.

```
[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
```

Unmount USB Device

To unmount the USB device in sysadmin-vm, use the following command:

usb device operation unmount

To unmount the USB device in XR-vm, use the following command:

unmount disk2:

Example

The following is an example of USB unmount in sysadmin-vm.

```
sysadmin-vm:0_RP0#usb device operation unmount
Fri Jul  13 09:25:24.531 UTC success usb unmounted
```

The following is an example of USB unmount in XR-vm.

```
RP/0/RP0/CPU0:ios#unmount disk2:
Fri Jul 13 14:56:46.393 IST
disk2: unmounted successfully.
```

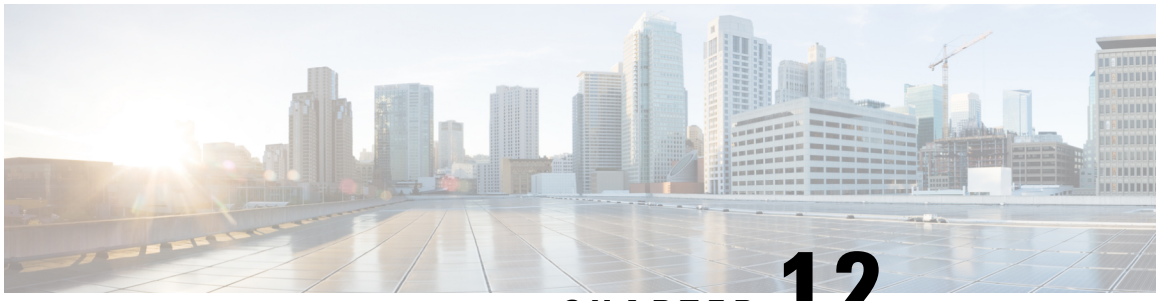

CHAPTER 12

Fault Profiles

This chapter describes how to configure and manage fault profiles.

- Fault Profiles, on page 125
- Tasks for Configuring Fault Profiles, on page 126
- Configure Fault Profiles, on page 126

Fault Profiles

The default fault list in a system captures all the possible type of faults that the system generates, along with the associated default severity values, for each fault type. This default severity value is the severity of the fault that is generated in a system when no other fault profile is defined and applied in that system. Based on your requirement, you can create new fault profiles and change the severity of fault.

The Fault Profiling feature enables you to create a unique fault profile for faults on the system or the line card. Each fault profile can contain one or more faults with user-defined severities. The highest precedence is maintained at the port level and the lowest precedence is maintained at the system level. For example, if the system profile is already attached and if you want to have a separate fault profile for a node, you can create a node profile and attach it to that node. The node inherits the properties of the node profile. The available severity levels are:

- 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

Limitations of Fault Profiles

The following are the limitations for fault profiles on Cisco NCS 1004:

- Fault profiling is available only on data path alarms—Optics, Coherent DSP, Ethernet, and ODU alarms.
- Fault profiling at the port level is not supported.
- You can create a maximum of 61 profiles.

Tasks for Configuring Fault Profiles

The following are the tasks for creating and configuring fault profiles:

- Create a fault profile with a unique name and a fault type.
- Add alarm names and severity level.
- Apply the fault profile at system-level or node-level.

Configure Fault Profiles

This task has details of how to create a fault profile and apply the fault profile at the system or node level.

configure

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*

commit

fault-profile *fault-name* **apply rack** *rack_id* **slot** { **ALL** | **LC** }

commit

exit

Examples

The following sample creates a fault profile and applies at system level.

```
RP/0/RP0/CPU0:ios(config)fault-profile FpSystem fault-identifier subsystem XR fault-type
HW_OPTICS fault-tag OPTICAL_LO_RXPOWER sas NONFAULTED nsas NONFAULTED
RP/0/RP0/CPU0:ios(config)commit
RP/0/RP0/CPU0:ios(config)fault-profile FpSystem apply rack 0 slot ALL
RP/0/RP0/CPU0:ios(config)commit
```

The following sample creates a fault profile and applies at node level.

```
RP/0/RP0/CPU0:ios(config)fault-profile FpNode fault-identifier subsystem XR fault-type
HW_OPTICS fault-tag OPTICAL_LO_RXPOWER sas CRITICAL nsas CRITICAL
RP/0/RP0/CPU0:ios(config)commit
RP/0/RP0/CPU0:ios(config)fault-profile FpNode apply rack 0 slot LC1
RP/0/RP0/CPU0:ios(config)commit
```

The following sample creates a fault profile, configures one second PM and applies at propagation level.

```
RP/0/RP0/CPU0:ios(config)fault-profile OTNAlarm fault-identifier subsystem XR fault-type
HW_G709 fault-tag G709_LOS sas CRITICAL nsas CRITICAL
RP/0/RP0/CPU0:ios(config)commit
RP/0/RP0/CPU0:ios(config)fault-profile OTNAlarm apply rack 0 slot ALL propagate
RP/0/RP0/CPU0:ios(config)commit
```

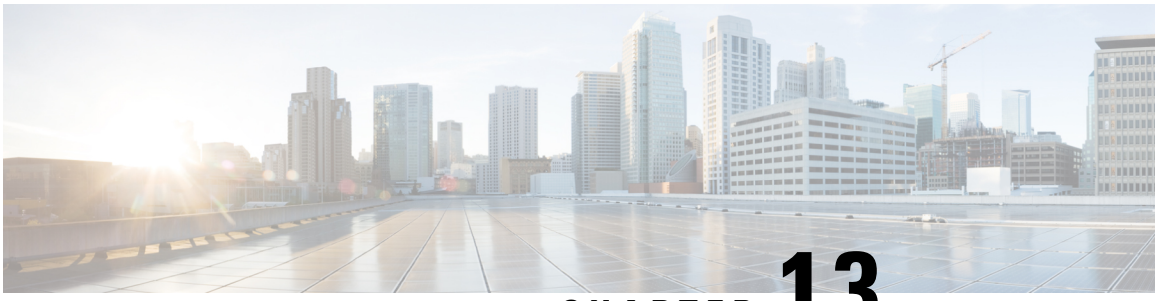

CHAPTER 13

Implementing Host Services and Applications

Cisco IOS XR software Host Services and Applications features on the router are used primarily for checking network connectivity and the route a packet follows to reach a destination, mapping a hostname to an IP address or an IP address to a hostname, and transferring files between routers and UNIX workstations.

Prerequisites for implementing Host Services and Applications

Ensure to install the relevant optional RPM package before using the host services or applications.

- HTTP Client Application, on page 129

HTTP Client Application

HTTP Client allows files to be transferred from http server to another device over a network using HTTP protocol. You can configure http client and various parameters associated with it by using the **http client** command.

Configure HTTP Client

HTTP Client application is available by default. You can configure http client settings or view and modify the existing settings. To configure the settings, use the **http client** command in XR config mode.

```
Router #configure
Router(config)#http client ?
connection          Configure HTTP Client connection
response            How long HTTP Client waits for a response from the server
                    for a request message before giving up
secure-verify-host  Verify that if server certificate is for the server it is known as
secure-verify-peer  Verify authenticity of the peer's certificate
source-interface    Specify interface for source address
ssl                 SSL configuration to be used for HTTPS requests
tcp-window-scale    Set tcp window-scale factor for High Latency links
version             HTTP Version to be used in HTTP requests
vrf                 Name of vrf
```

Table 11: Commands used to configure HTTP Client settings

Features	Description
connection	Configure HTTP Client connection by using either retry or timeout options.

Features	Description
response	How long HTTP Client waits for a response from the server for a request message before giving up.
secure-verify-host	Verify host in peer's certificate. To disable verifying this, you can use the command **http client secure-verify-host disable**
secure-verify-peer	Verify authenticity of the peer's certificate.
source-interface	Specifies the interface for source address for all outgoing HTTP connections. You can enter either an ipv4 or ipv6 address or both.
ssl version	SSL version (configuration) to be used for HTTPS requests.
tcp-window-scale scale	Set tcp window-scale factor for high latency links.
version version	HTTP version to be used in HTTP requests. • 1.0 - HTTP1.0 will be used for all HTTP requests. • 1.1 - HTTP1.1 will be used for all HTTP requests. • default libcurl - will use HTTP version automatically.
vrf name	Name of vrf.

Examples

Example 1: This example shows how to set the tcp window-scale to 8.

```
Router(config)#http client tcp-window-scale 8
```

Example 2: This example shows how to set the HTTP version to 1.0.

```
Router(config)#http client version 1.0
```

Note HTTP Client uses libcurl version 7.30

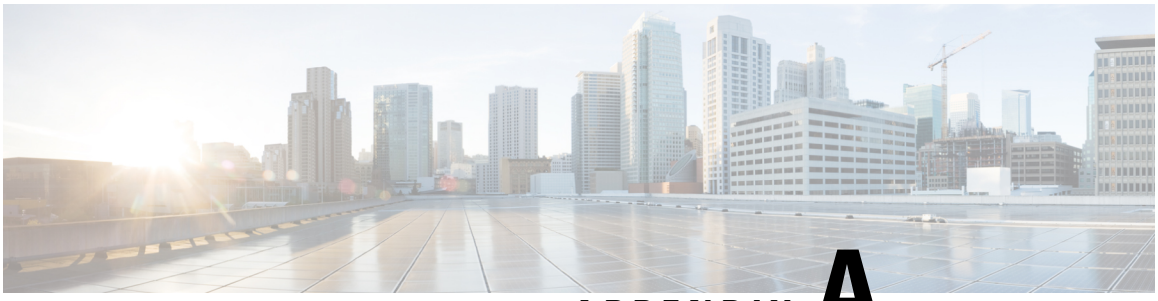

APPENDIX A

SNMP

NCS 1004 supports the following MIBs.

- 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
- 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.